



VOLUME: 07 ISSUE: 3 May 2009 116 PAGES ISSUE# 76

For You

Operating Systems that can set you free

PC-BSD, PCLinuxOS, Pardus, Dreamlinux & SimplyMEPIS reviewed. Want to try them? Check out the discs.

Accidentally
Deleted Files?
Recover Easily | 46

How to
Configure Tablet
Digital Pen | 44

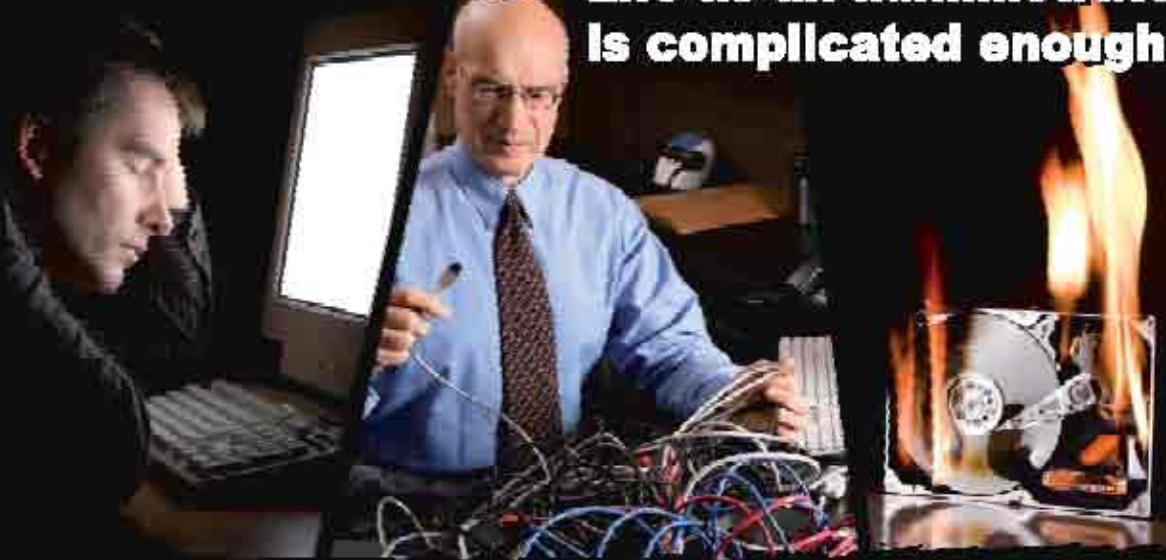
Automate Work
with OOo Macros | 52

Contain
Linux Instances
with OpenVZ | 66

Balance Traffic
Across Data Centres with LVS | 76

Develop
Interactive Web Pages
with JQuery | 92

**Life as an administrator
Is complicated enough...**



**But Backup and Recovery for your
Linux Servers does not need to be.**

NetVault: Backup simplifies backup and recovery
without compromising **Functionality and Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

Permanent FREE Use Edition for Linux is available for download at <http://www.bakbone.com/nvbu/redhat/>



NetVault is a true Linux data protection solution featuring:

- Online backup
- Point and click recovery
- Fully automated protection and recovery support for Linux based applications
- Virtual Tape Library (VTL) with disk staging
- SAN support with LAN free backups
- Backup duplication for off-site storage of backups
- Support for MySQL native replication, restore to table level and to alternative databases or instances.

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 36 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

- LLC102: Essentials of Linux OS
- LLC103: Linux System & Network Administration
- LLC203: Linux Advanced Administration
- LLC303: Linux System & Network Monitoring Tools
- LLC403: Qmail Server Administration
- LLC404: Postfix Server Administration
- LLC405: Linux Firewall Solutions
- LLC406: OpenLDAP Server Administration
- LLC408: Samba Server Administration
- LLC409: DNS Administration
- LLC410: Nagios - System & Network Monitoring Software

Courses for Developers

- LLC104: Linux Internals & Programming Essentials
- LLC106: Programming with Qt
- LLC106: Device Driver Programming on Linux
- LLC107: Network Programming on Linux
- LLC108: Bash Shell Scripting Essentials
- LLC109: CVS on Linux
- LLC204: MySQL on Linux
- LLC206: Programming with PHP
- LLC206: Programming with Perl
- LLC207: Programming with Python
- LLC208: PostgreSQL on Linux
- LLC501: Programming with OpenGL
- LLC504: Linux on Embedded Systems

RHCE Certification Training

- RH038: Red Hat Linux Essentials
- RH133: Red Hat Linux System Administration
- RH253: Red Hat Linux Networking & Security Administration
- RH300/301: Red Hat Rapid Track Certification Course

RHCSA Certification Training

- RHS233: Red Hat Enterprise Security: Network Services
- RH423: Red Hat Enterprise Directory Services & Authentication
- RHS428: Red Hat Enterprise SELinux Policy Administration

**4 Day Full Time Course on
Postfix Administration
starting on 9 May 2009**

**4 Day Fast Track Course on
Linux on Embedded Systems
starting on 9 May 2009**

**RHCE Exams in Bangalore
on 15 & 29 May 2009**

**RH301 from 11 & 25 May '09
RHS333 from 9 May '09
RH423 from 16 May '09
RHS429 from 23 May '09**

**RHCSA Exams in Bangalore
on 23, 24 & 25 June 2009**

Linux Support & Solutions

Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS Servers

For more info log on to:
www.linuxlearningcentre.com
Call: 9845057731 / 9343780054
Email: info@linuxlearningcentre.com

RED HAT Training Partner
RHCE & RHCSA Exam Centre



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

#635, 6th Main Road, (Ad.- Bank of India) Hennurmathnagar, Bangalore 560019

Tel: +91.80.22428838 / 26600839 / 26610999 Tel/Fax: +91.80.26808839

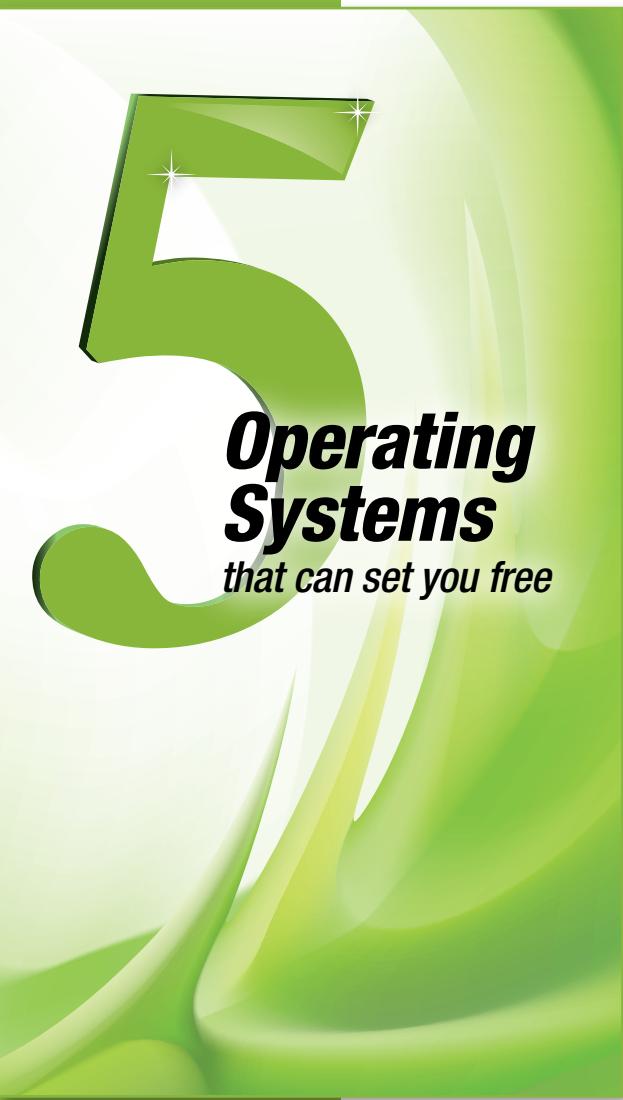
Call: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com
www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore

1291, 24th Cross, 30th Main,

BSK II Stage, Bangalore 560070

Tel: +91.80.26712928



FOR YOU & ME

- 20 Laid Off? Why Apply For a New Job When You Can Be a CEO?
- 22 Is It a Bird? Is It a Plane? No! It's BSD!
- 26 Still Radically Simple
- 30 Kaptan At Your Service
- 34 Is It A Dream or What?
- 38 A Distro Hopper's Final Destination?
- 42 Run Linux on Windows, Part 2
- 44 An Artist's Haven
- 46 Recovering Deleted Files Easily



GEEKS

- 52 Automate Your Work with OpenOffice.org Macros
- 56 Fast, Efficient and Reliable Pattern Scanning and Processing
- 60 Building a Server from Scratch—Part 4: Pages Full of Data
- 84 Programming in Python for Friends and Relations—Part 13: Creating Rhythmic Noise



ADMIN

- 64 Open Source for Career Management
- 66 Containing Linux Instances with OpenVZ
- 72 The Art of Guard, Part 2
- 76 Balancing Traffic Across Data Centres Using LVS



CONTENTS

DEVELOPERS

- 88 Testing RESTful WebServices Made Easy
- 92 Develop Interactive Web Pages with Ease
- 96 Subversion 1.6: Security Improvements Illustrated
- 98 Database Programming in Python

COLUMNS

- 41 FreedomYug: Doc Till You Drop
- 50 Foss on Mobiles: Waiting for the Android Effect!
- 82 CodeSport
- 87 The Joy of Programming: Silly Programming Mistakes => Serious Harm!
- 102 A Voyage to the Kernel—Part 12: Segment 3.1, Day 11



REGULAR FEATURES

- 06 Editorial
- 08 You Said It...
- 10 Technology News
- 14 Industry News
- 18 Q&A Section
- 80 Tips & Tricks
- 91 DVD Page
- 106 Linux Jobs
- 108 FOSS Yellow Pages



LFY DVD: PC-BSD 7.1 & more

PC-BSD is a free operating system with usability in mind. Like any modern system, you can listen to your favourite music, watch movies, work with office documents and install additional applications using setup wizards just like in Windows.



LFY CD: PCLinuxOS 2009.1

This release features kernel 2.6.26.8.tex3, KDE 3.5.10, OpenOffice.org 3.0, Firefox 3.0.7, Thunderbird 2.0.0.14, Ktorrent, Frostwire, Amarok, Flash, Java JRE, Compiz-Fusion 3D and much more.

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgrove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares
D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA
D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI
Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE
Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD
P.S. Muralidharan
Ph: 09849962660
E-mail: efhyd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD
Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 24942538, 24925651,
24927383 Fax: 24950392
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the license. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear Readers,

Oracle has taken over Sun Microsystems.

Can't think of any other development that could beat this news, in terms of its sheer impact on the Open Source world. This one is really big! And with so many possible eventualities, many of which will give open source enthusiasts the goose bumps. We decided to do a major story on this development for our next issue (June 09), in order to explore every possible implication of this mega merger.

The question that topped our list was what would happen to MySQL, which Sun had recently acquired for a whopping US\$ 1 billion? Many had feared for MySQL's long-term prospects after Sun's acquisition—but the lure of reaching out to enterprise customers through Sun's sales team seemed to offset the threats. Now, we are talking about the biggest proprietary database firm acquiring the fastest growing open source database project—that's tricky. We did ask a few experts for their opinion on the subject, and got mixed views. Some believe that MySQL is now too strong to be killed, and hence Oracle would be smart enough to use it and expand its base in new markets, while others believe it might try to slow down MySQL's growth so that it does not end up competing with Oracle's own database.

What about OpenOffice.org? Will Oracle use it to create problems for Microsoft on MS Office turf? After all, Larry Ellison, Oracle's big boss, has been quite open about his anti-Microsoft sentiments, and this could be a good chance for him. Plus, OpenOffice.org also powers the Symphony project of another of Oracle's competitors—IBM. How closely will the two teams work, I wonder.

Red Hat, too, could have some serious competition now. Earlier, Oracle was selling Red Hat's Enterprise Linux (RHEL) minus the Red Hat trademarks. Now Oracle has access to the entire Solaris stack, and if it combines the strength of Solaris and Linux, and leverages its sales army—we could have a serious battle. Noting that this will only

result in even more options for customers, Jim Zemlin of the Linux Foundation has gone public in stating that it's 'good news' for the Linux ecosystem.

Then there's Java too, which was also open sourced by Sun. Ellison is beaming with pride about this acquisition. His famous quote about Java says it all, "...the most important software Oracle has ever acquired". So, there's nothing to worry about here. But, if you are connected with other Sun-sponsored Open Source projects such as NetBeans and GlassFish, there're reasons for concern. Almost everyone believes that Oracle will slash Sun's R&D budget, and if that happens, these projects might get disbanded.

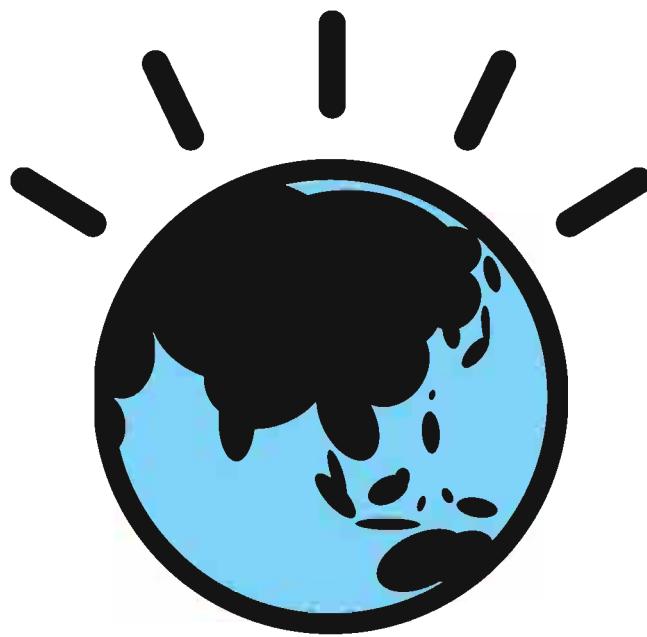
MySQL is now too strong to be killed, and hence Oracle would be smart enough to use it and expand its base in new markets

In a nutshell—Linux has seemingly got a boost. Java's future is certain. The fate of NetBeans and GlassFish is threatened, and that of MySQL and OpenOffice.org is anyone's guess.

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com



A mandate for change is a mandate for smart.

The mandate for change today is not merely for political leaders, but also for managers of businesses everywhere.

These volatile times have put the onus on businesses to focus on complex global systems now more than ever. And we have had an important learning from engaging in these systems: global integration impacts the way the world works.

We have also realized, as the world gets 'smaller' and 'flatter', merely being connected is not enough to survive, let alone thrive.

Fortunately, something is happening right now: our planet is becoming 'smarter'.

Here is how. With billions of inexpensive sensors being embedded inside everything, from trains and planes to livestock and medicines, our world is becoming instrumented, allowing systems and objects to 'speak' with each other. Linked to powerful backend systems, these instruments can analyze data and turn information into insights in real time, turning mere assets into intelligent assets.

With so much ready potential, what would you not change?

Consider these:

The second annual Global Retail Theft Barometer Survey across 36 countries in 2008 found that India has the highest shrinkage rate at 3.1% (that's about Rs. 12,392 crores in losses).

Over 87% of India's poorest households have no access to credit at all.

Owing to its unique geo-climatic conditions, India is highly prone to natural disasters. So much so, they cost the country 13% of its GDP today and will go on to become a major stumbling block to its economic growth by next year.

The good news is, on a smarter planet, these problems are solvable.

Now, consider these:

Hindustan Petroleum Corporation Ltd. is now in the process of implementing a tracking system for its Liquefied Petroleum Gas cylinders from the bottling plant to its distributor network, using a Radio Frequency Identification (RFID) based solution, helping curb illegal diversions, while delivering an essential resource to consumers.

Financial Information Network and Operations Ltd., a technology solutions provider focusing on micro customers, is helping many of the nation's disadvantaged gain access to simple financial products to improve their quality of life.

While it cannot eradicate natural disasters or even reduce them, technology can certainly help mitigate losses and alleviate human suffering. Based on analysis of affected regions, relief agencies can be guided to manage people, resources and logistics in a far more efficient manner.

Clearly, there is a pressing mandate for change today. And we have the resources to get started. Log on to ibm.com/think/in and be part of the revolution with IBM.

You said it...



I have a few comments on the new features of SVN 1.6.

I'd be really interested to know if SVN now tracks the file history across a simple rename—something that is likely to be of far more universal interest than an actual conflict. In any case, Git (and possibly other VCSs) seem to do just the right thing even when there is a conflict rename/modify conflict such as the one described, and usually without user input.

On the security aspect, authentication is no longer something that a normal 'application' should deal with; delegating it to tried and tested mechanisms like SSH and public key authentication is the only sensible option. The question is not how the application stores the password but why the application should even have to worry about this. Not all applications can piggyback on SSH, but I see no reason why SVN could not have done so.

Finally, I also wonder why it is not possible to unpack a 'packed' repository—one-way operations in a code repository are somewhat disconcerting, to say the least. Is it merely that no one has written an unpacker (which is OK) or that it is actually impossible (which sounds like the pack format is fragile in some way)? I certainly hope it is the former.

I certainly don't agree with Linus Torvalds' (in)famous statement that, "SVN is the most pointless project ever started." I regard SVN as walking the minefield, so the ones that follow know what to avoid, and that is worthy of our respect and gratitude.

But I also feel (and I have said this before) that continuing on a project that started out as "CVS done right" is not an optimal use of the enormous talent and brains that are behind SVN. CVS is more than 18 years old now, and the time for fixing its problems is long gone. The right thing to do is to give it an honourable retirement and let the new crop take over.

—**Sitaram Chamarty, sitaramc@gmail.com**



I read the article about 'FOSSconf 2009' in the April 2009. I wish to share a few more things that happened:

1. Most FOSS lovers visited the conference without thinking about time and distance. A few people came from Bangalore, Hyderabad, etc. One man from Canada came to this event and visited all demo stalls, enquiring a lot about the localisation of Linux.
2. Students from TVS Matriculation, in the age group of around 12 years, really rocked the demo stalls.
3. The Chennai LUG members issued cash gifts of Rs 10,000 each to the best demo stall, best project and best new student speaker to encourage the students in FOSS activities.

I was there at the demo stall of the Ubuntu Tamil team.

—**Padhu, Pollachi**



The April issue was very good, especially the article on how Dell could eat up Apple's market share. As for Debian Lenny, I am dying to get the other four DVDs. I've started reading your magazine since last year and I'm in love with the magazine now. The articles on Python are really good.

A small request: please include an article on Network Simulator 2 and 3. It would be very helpful for network designers like me. Also, I would really appreciate it if the magazine bundled distros like Linux Mint and Lindows in the forthcoming issues. And, by the way, keep up the good work!

—**Nitish, nitishbanger@yahoo.co.in**



ED: Wow! Thanks for your compliments. We'll certainly try our best to include articles on network simulators in the upcoming issues. As for your request to bundle Linux Mint and Lindows (I guess you meant Linspire, which is now Freespire?) I've conveyed your message to the CD team.



I was really upset after seeing www.linbai.info/magazine/linux-for-you-magazine-march-2009.html. Note that the link might change but the soft copy of *LFY* is available for download in that website, which was previously called www.ebookpedia.info.

I hate to see my all time favourites, like *LFY*, coming out as e-books, scanned by some reader. *LFY* folks—please do the needful. Is the e-book publishing plan required for *LFY*?

—**Gouri Ananth, ananth.gouri@gmail.com**

ED: Well, this is certainly interesting. There's nothing much we can do as most of the content in the magazine is released under the Creative Commons-Attribution-ShareAlike 3.0 Unported License after a month from the date of publication. Anyway, work is on at LinuxForU.com and the site should be up by the end of May, if things go as per plan. Once it goes live, we'll slowly start feeding in the magazine archives. Please don't expect every issue at once—May '09 will be our 76th issue. Anyway, we'll certainly welcome all those who're interested in reading the PDFs, to come and check out all our archives. Our licensing permits everyone to use most of our articles in any way, as long as they give attribution, and release the modified content under the same licence. :-) We only hope, a Web interface makes this easier.

Errata

In our April '09 issue, on the Contents page, the description of OpenSolaris was wrongly repeated against Debian Lenny too. We regret this error.

Please send your comments or suggestions to:

The Editor



LINUX FOR YOU Magazine
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: 011-26810601/02/03
Fax: 011-26817563
Email: lfyedit@efyindia.com
Website: www.openTis.com

Call for participation NCOSS-09

National Conference on Open Source Software
May 25th - 26th, 2009, Navi Mumbai, India.



Organized by Centre for Development of Advanced Computing, Mumbai
Supported by IEEE Computer Society, Mumbai and Chennai chapters & CSI, Div II on Software & SIG-OSS
Media partner Linux for you

The National Conference on Open Source Software is intended to act as a forum for promoting adoption of open source softwares and solutions in different areas, by sharing experiences in solution selection, customisation/adaptation, etc. Over 110 papers were submitted to the conference from all over India. After a rigorous international refereeing process, about 30 papers have been selected, which will be presented during the conference.

Conference Programme

- Paper Presentation in areas including
 - e-Learning
 - Collaboration Technology
 - Language Computing and Localisation
 - Knowledge Management
 - e-Health
 - Machine Learning and Data Mining
 - Testing and Application Frameworks
- Invited Talks
- Panel Discussions
- Live demos of FOSS applications
- Pre-Conference Tutorials (on 24th May)

Demo Hall

NCOSS-09, announces DEMO HALL for showcasing open source applications. The aim is to provide a platform for people to demonstrate their open source works/tools. Irrespective of whether you have a paper at the conference, you can use this opportunity to showcase your work.

All those participating in the demo hall must register for the conference. Mail your applications to ncoos@cdacmumbai.in giving a brief description of the demo sessions by 5th May 09. For more details visit <http://ncoos.cdacmumbai.in>

Registration Fees

	Conference	Tutorial
Not for profit/ educational institutes	Rs. 2200	Rs. 700
Industry	Rs. 3500	Rs. 1000
Student	Rs. 1100	Rs. 500

(Registration form and more details on website)

Programme Committee

Ramakrishnan S., C-DAC, Pune (chair)
Chakravarti A. K. IIT, Delhi
Zia Sequeib, C-DAC, Mumbai
Sasikumar M, C-DAC, Mumbai
Sivakumar G, IIT, Mumbai
Venkatesh Hariharan, RedHat, Mumbai
Mohan H R, The Hindu, Chennai
Rishab Aiyer Gosh, UNU-MERIT, Netherlands
Krishnan C N, AU-KBC, Chennai
Gora Mohanty, Samsi, Delhi
Rajagopal M R, C-DAC, Chennai
Dravida Seetharam, IBM, Bangalore
Nagarjuna G, HBCSE, Mumbai

Contact:

 (For more details and registration)

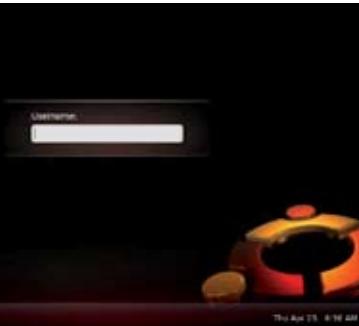
NCOSS-09 Secretariat
C-DAC Mumbai, Raintree Marg, Near Bharati Vidyapeeth,
Opp. Kharghar Railway Station, Sector 7, CBD Belapur,
Navi Mumbai - 400 614, INDIA.
E-mail: ncoos@cdacmumbai.in
Phone: +91-22-27565303, Fax: +91-22-27560004
URL: <http://ncoos.cdacmumbai.in>

Technology News

Jaunty Jackalope's out! Oh, that's the latest Ubuntu

Canonical, the commercial sponsor of Ubuntu, has launched the Ubuntu 9.04 Desktop and Server Netbook Remix editions.

Ubuntu 9.04 Desktop Edition delivers shorter boot speeds, enhanced suspend-and-resume features, and intelligent switching between Wi-Fi and 3G environments. It includes the latest GNOME 2.26 desktop environment with a number of great



new features, including Brasero 2.26.0, an all-in-one CD burning application along with the default disc burning utility in Nautilus. It also comes with improved handling of multiple monitors; X.Org server 1.6; Wacom tablet hotplugging; a new style for notifications and notification preferences; significantly improved boot performance; Linux kernel 2.6.28; and optional ext4 file system support.

Meanwhile, Kubuntu 9.04 includes KDE 4.2 with many new features like significant refinements of Plasma and KWin, the KDE workspace, many new and updated Plasma widgets, new and improved desktop effects (enabled by default), the return of the optional 'Classic Desktop' motif as an option, etc. The new inclusions in System Settings are tools for managing software and printer configuration; Quassel, a new IRC client; Amarok 2.0.2, KTorrent 3.2, Digikam 0.10.0, etc.

Enhancements in the Server Edition include improved virtualisation with the latest KVM features, clustering support in the Samba file server, and easier mail server set-ups with out-of-the-box Dovecot-Postfix integration.

Ubuntu 9.04 Netbook Remix comes with faster boot speeds, and with a built-for-purpose interface, which means that favourite applications and websites are just a click away. For the first time, users can download the complete Ubuntu Netbook Remix to a USB flash drive directly from Ubuntu.com. Users can then install and run Ubuntu Netbook Remix on a wide range of the most popular netbook machines available in the market today. Ubuntu 9.04 Netbook Remix has been fully tested for use on a range of netbook models, including Acer Aspire One, Asus eee PC 1000 and Dell Mini 9.

ISC kickstarts development of BIND 10

The Internet Systems Consortium (ISC), with the support of some sponsors, has revealed plans for BIND 10, the next leap forward in DNS server software. BIND 10 promises to be modular, highly scalable, and provides simple methods for configuration management and integration with other systems. The design goals are simple: a secure, flexible, resilient DNS server that integrates easily into the workflow and in the maintenance that complex networks demand. Additionally, BIND 10 has set a goal to provide the state-of-the-art in DNS security. The differentiation will be the way that a user configures the secure services they choose to deploy. The design goal for DNSSEC in BIND 10 is to be usable by the typical DNS administrator with built-in safeguards for management and renewal. For more information on the project and about participation, visit www.isc.org/bind10.



Development kit enables digital photo frames programming

RMI Corporation, a fabless semiconductor company that provides high-performance processors for communication and media-rich applications, has launched its Home Media Player Application Development Kit (ADK). The ADK is an open development system and includes a complete software Board Support Package (BSP) for users and original design manufacturers (ODMs) targeting media-based applications, such as wireless Digital Photo Frames (DPFs), Mobile Internet Devices (MIDs), Digital Signage, Information Dashboards and other home networking and consumer devices.

The hardware includes a 10.4-inch (26.4 cm) high resolution (1024x768) LCD, 2 GB of NAND flash, 256 MB of DDR2 RAM, Wi-Fi connectivity (through partnership with Atheros), two USB ports, Ethernet and serial ports, stereo speakers, IR receiver, and a 5-in-1 card reader. The ADK utilises the Linux OpenEmbedded environment to create a Board Support Package based on Linux 2.6.x. It

includes a media stack and player from CoreCodec for D1 resolution video playback (MPEG-1/2/4, VC1/WMV, H.264, DivX, XviD), hardware-accelerated DirectFB, GTK+ middleware, a complete tool chain, filesystem, and an example application.

The ADK is available directly from RMI or through its distribution channels for \$499.

No gatherings or get-togethers, yet 6 billion conversations everyday.

Making the impossible possible, is a way of life at Yahoo! Small or big, whatever be the idea, we bring it to life. No wonder, Yahoo! has constantly conceived and delivered world-class products that touch the lives of over 500 million users. Whether it's email, news & entertainment, sports, music or just about anything else on the net, people turn to Yahoo!. Making it the world's most visited site and among the most well-known brands globally. Yahoo! India R&D is a strategic Center that contributes in making Yahoo! what it is.

We are presently looking for expert Linux programmers with BTech/MTech/MS/PhD in Computer Science for the following positions:

Architect, Principal Engineer - Cloud Computing

The Cloud Computing group at Yahoo! Bangalore is chartered with making our data center environment friendly and efficient. We are looking for an Architect and Principal Engineers.

The ideal candidate should have: • M.Tech/MS in Computer Science with 9+ years experience • Strong architectural and/or debugging and troubleshooting skills in C++/Java on UNIX. in one or more of the following areas below • Database - 6+ years of experience in DB schema design, PL/SQL, DB performance tuning for databases like MySQL, Oracle & PostGres • Distributed and Parallel computing - 5+ years of experience in designing complex enterprise /internet scale multi-tier/P2P/multithreaded applications addressing requirements like high performance, fault tolerance, maintainability & serviceability • UI & UI Framework - 2+ years of experience in defining REST/SOAP web-services. 4+ years of experience in creating robust UI frameworks that are also scalable.

Performance Architect - Performance Engineering

The Performance Engineering team is an expert group of systems programmers that analyzes and improves

performance and scalability of Yahoo! Search, Advertiser-Publisher Exchange and Grid Technology. This group is currently building presence in Bangalore and is hiring Architects and Principal Engineers.

The ideal candidate should have: • M.Tech/MS/PhD in Computer Science with 9-16 years experience • Deep knowledge of Unix/Linux OS internals or JVM internals • Proficiency in Systems Programming using C/C++ or Java • Knowledge of Computer Architecture of modern Intel/AMD processor/chipsets • Background in High Performance Computing (HPC) or Performance Optimizations is a plus • Knowledge of Apache Hadoop internals (<http://hadoop.apache.org>) is a plus

Tech Lead - Linux Engineering

The yLinux team is the "Linux Center of Excellence" for all of Yahoo!. This team is responsible for Linux strategy, support and development. This team works closely with Yahoo! properties, Linux vendors and the open source community to provide a stable, secure and scalable Linux infrastructure to Yahoo!. The team comprises of kernel, systems and release engineers that provide solutions spanning Linux Kernel support, migration of code from FreeBSD to Linux and packaging/public domain software.

The ideal candidate should have: • B.Tech in Computer Science with 7+ years experience. M.Tech/ MS is a plus • Extensive development experience on a Linux environment • Good understanding of building, packaging software on Linux, including RPMs • Excellent C/C++ development and debugging skills • Background in porting source code to Linux is a plus • Scripting skills in Perl/PHP & Unix shell programming is a plus

Positions are based in Bangalore.
Mail in your resume to career@yahoo-inc.com

YAHOO! RESEARCH & DEVELOPMENT
INDIA

For more information log to:
<http://india.yahoo.com>

• 500 million users • 4 billions clicks every day • 16 languages • 150+ services

Technology News

VirtualBox 2.2 supports OVF and virtual appliances

Sun Microsystems has announced the availability of Sun VirtualBox 2.2, which introduces support for the new Open Virtualization Format (OVF) standard, along with significant performance enhancements and updates. VirtualBox 2.2 software enables users to build virtual machines or appliances and effortlessly export them from a development environment, and import them into a production environment. Support for OVF also helps to ensure VirtualBox 2.2 software is interoperable with other technologies that follow the standard.

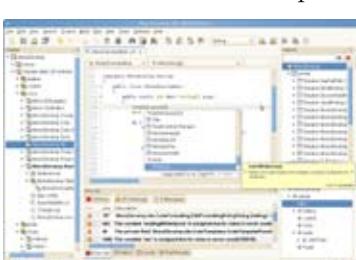
Additional features of VirtualBox 2.2 software include Hypervisor optimisations, 3D graphics acceleration for Linux and Solaris applications using OpenGL, and a new host-interface networking mode, which makes it easier to run server applications in virtual machines. To download the freely available Sun VirtualBox software, visit: <http://www.sun.com/software/products/virtualbox/get.jsp>.



Better ASP.NET with MonoDevelop 2.0, Mono 2.4

MonoDevelop 2.0 has been released with the aim of enabling developers to write desktop and ASP.NET Web applications on Linux, port .NET applications created with Microsoft Visual Studio to Linux and Mac OS X, and maintain a single code base for all three platforms. MonoDevelop provides tools to simplify and

streamline .NET application development on Linux, including: improved ASP.NET and C# 3.0 support and a built-in debugger. It now uses MSBuild-style project files to increase interoperability with Visual Studio. Web projects are now also compatible with Visual Studio 2008 and Visual Web Developer 2008 SP1, providing more options for developers who want to build and deploy their Web applications on both Windows and Linux.



New features available in Mono 2.4 include a new code generation engine that improves the performance of executing .NET applications on the Mono runtime, while managed Single Instruction, Multiple Data (SIMD) extensions enable developers to take advantage of hardware acceleration without having to program in lower-level languages. Additional runtime innovations, such as full ahead-of-time (AOT) compilation, bring Mono-based applications to new platforms, including the Apple iPhone. MonoDevelop 2.0 and Mono 2.4 are available now and can be downloaded at www.mono-project.com/downloads.

New Embedded InnoDB

Innodb has introduced its newest product, Embedded InnoDB. Made for application developers, device makers and ISVs, Embedded InnoDB claims to provide all the high-performance, proven and reliable data management features of InnoDB for embedding in target applications—all with a small footprint, low overhead and no unnecessary features. It is designed to be linked directly into application programs, and provides highly efficient, low-level database management services, without using SQL. Embedded InnoDB is currently available in binary form for Linux and Windows, and in source form under the GPL V2 licence.

The Android 1.5 early-look SDK is out!

To give developers a headstart, the next release of the Android 1.5 early-look SDK has now been made available. The Android 1.5 platform will include many improvements and new features for users and developers. Additionally, the SDK itself introduces several new capabilities that enable you to develop applications more efficiently for multiple platform versions and locales. The new SDK has a different component structure compared to earlier SDK releases. This means that it does not work with older Eclipse plug-ins (ADT 0.8) and the old SDKs (1.1 and earlier) do not work with this new Eclipse plug-in (ADT 0.9).



Note that since this is an early-look SDK, the tools and documentation are not complete. Additionally, the API reference documentation for it is provided only in the downloadable SDK package—see the documentation in the SDK's *docs/reference/* directory. Visit developer.android.com/sdk/preview/features.html to get a grip over the feature set of the Android 1.5 SDK.

CentOS-5.3 for i386 and x86-64

Based on the upstream release of RHEL 5.3, CentOS 5.3 has finally been released and introduces a completely new artwork stack. The contribs repository is also now back. The release notes for CentOS-5.3 can be found on-line at wiki.centos.org/Manuals/ReleaseNotes/CentOS5.3.

WEB INNOVATION 2009

WEBSERVICES / MOBILE / MICROFORMATS / LONG TAIL / INDICATION / OPENAPI / XML STANDARDIZATION

SIMPLIFICATION / XML / SEARCH / JOY / USE / LIST / PREDICTION / USABILITY / AJAX / PAY PER CLICK / CSS

DESIGN / AGILE / OPS / PAY / ANALYTICS / SEO / HAVING / VOID / STRING / LISTS / DESIGN / RSS / ECONOMY

MOBILITY / OFC / THE LOM / SYNC / SVG / OPENAPI / XML STANDARDIZATION

BLOGS / AGILITY / XHTML / MANTICORE / OF USE / TRUST / PARTITION / USABILITY / AJAX / PAY PER CLICK / CSS

VIDEOCASTING / OPENDATA / RANK / VIDEO / SEO / SHARING / PODCASTING / LISTS / DESIGN / RSS / ECONOMY

WEB INNOVATION 2009

The Nextgen Web Technology Revolution: 2.0 & Beyond

New Delhi • Mumbai • Bangalore

IT'S TIME AGAIN TO CELEBRATE THE SIZE, POWER, AND INNOVATION OF THE WEB INDUSTRY

Web 2.0 has provided some of the most innovative technologies and methodologies since the creation of the Web itself. Blogs and wikis are becoming commonplace in the enterprise, and social networks are starting to gain traction. Add mashups, Ajax and RIA technologies, infuse consumerization, and you've got a whole new world of support and management challenges.

Some of the key touch points at these multi-city Conferences include:

- Understanding new technologies and methodologies associated with Web 2.0 to effectively utilize them
- Learning best practices regarding the Web from the experiences of leaders and innovators
- Identifying which aspects of the Web are ready for prime time, and which are too risky due to immaturity
- Discussing with the vendors providing Web Services and capabilities and identify the strongest players, vendors and newer web markets
- Connecting with the real-world solutions and strategies of your peers through the online community of registered attendees onsite networking, end-user case studies, panel discussions and more...

As Web 2.0 grows up and gets serious, the time is right to call upon the Indian WEB COMMUNITIES again and take this bandwagon add other Indian cities under the umbrella of the Web Innovation series. It's time again to come together - to learn, share and network. And celebrate the size, power, and innovation of the web industry. Web Innovation 2009 is for builders of the next generation web designers, developers, entrepreneurs, marketers, business strategists, and venture capitalists. WEB INNOVATION 2009 series will take the pulse of the Web ecosystem and looks to its future, and how the Internet Revolution is being created and delivered.

To register online, visit:
www.webinnovation.in/register.html

contact Naveen Lawrence +91 22 6702 2746
OR email naveen@tfci.com

REGISTER
TODAY

Conference Fee
Rs. 5,000/-
per delegate

April 29, 2009
The Leela Kempinski, Gurgaon
Delhi NCR

May 13, 2009,
Grand Hyatt
Mumbai

Media Partner



WHAT CAN I
DO WITH
THE POWER
OF WEB 2.0

HOW CAN I
USE THE POWER
OF WEB 2.0

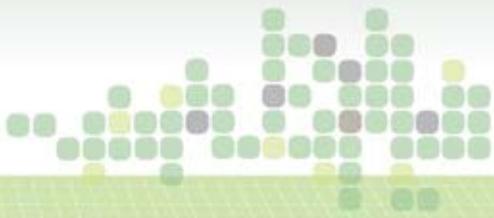
PROFITING
FROM
WEB 2.0

Supported by



Conceived, Produced & Organized by
 TRADE FAIRS &
CONFERENCES
INTERNATIONAL
New Frontiers In Business Marketing

Industry News



Oracle acquires Sun Microsystems

Oracle Corporation has agreed to buy California-headquartered Sun Microsystems in a \$7.4 billion (\$9.50 per share) deal. Sun's acquisition by Oracle caught many in the industry by surprise as the deal came just two weeks after IBM abandoned its bid to buy Sun. According to analysts, the deal strengthens Oracle's position against IBM and HP for dominance of the server and storage markets.



© Edie Avari

Kapil Dev Singh, country manager, IDC India, said, "The Oracle-Sun merger brings together two strong brands with strengths in the software and hardware space, with the combined entity aspiring to make a mark in the enterprise IT space. Sun's strong computing platform and Oracle's middleware and database platforms will make Oracle's IT infrastructure offering stronger."

There are substantial long-term customer advantages to Oracle owning two key Sun software assets: Java and Solaris. Java is one of the most widely deployed technologies, and it is the most important software Oracle has ever acquired. Oracle Fusion Middleware, Oracle's fastest growing business, is built on top of Java. Oracle can now ensure continued investment in Java technology for the benefit of customers and the Java community.

The Sun Solaris operating system is supposed to be the leading platform for the Oracle database, Oracle's largest business, and has been so for a long time. With the acquisition of Sun, Oracle can optimise the Oracle database for some of the high-end features of Solaris. Oracle says it is as committed as ever to Linux and other open platforms, and will continue to support and enhance strong industry partnerships.

"Oracle will be the only company that can engineer an integrated system where all the pieces fit and work together so customers do not have to do it themselves."

"We expect this acquisition to be accretive to Oracle's earnings by at least 15 cents on a non-GAAP basis in the first full year after closing. We estimate that the acquired business will contribute over \$1.5 billion to Oracle's non-GAAP operating profit in the first year, increasing to over \$2 billion in the second year. This would make the Sun acquisition more profitable in per-share contributions in the first year than we had planned for the acquisitions of BEA, PeopleSoft and Siebel combined," said Safra Catz, president, Oracle.

"The acquisition of Sun transforms the IT industry, combining best-in-class enterprise software and mission-critical computing systems," said Larry Ellison, CEO, Oracle. "Oracle will be the only company that can engineer an integrated system—applications to disk—where all the pieces fit and work together so customers do not have to do it themselves. Our customers benefit as their systems integration costs go down, while systems performance, reliability and security, go up."

Ingres joins Open Source Channel Alliance

Ingres Corporation, an open source database management company, has joined the Open Source Channel Alliance, a group founded by Red Hat and SYNEX Corporation, to extend the value of open source to a broad set of customers through the 15,000 resellers of SYNEX. Ingres joins as a charter member and has signed a distribution agreement with SYNEX.

The alliance enables SYNEX to distribute the Ingres product portfolio as well as other open source applications, as part of an end-to-end open source solution approach that will deliver a high return-on-investment (ROI) as IT budgets continue to shrink. Specifically, Ingres Database, an open source database that helps organisations develop and manage business critical applications at an affordable cost, will be available through SYNEX' distribution channel.

Azingo joins OMTP

Mobile Linux company, Azingo, has joined the Open Mobile Terminal Platform (OMTP) to advance the BONDI initiative, which ensures that developers can create secure mobile Web applications and services across different platforms and mobile devices. Azingo has developed Azingo Mobile 2.0, an open mobile platform, to enable handset makers to deliver next-generation Web 2.0 experiences on mobile phones. By joining OMTP, Azingo said it plans to contribute its expertise in open mobile platform development to the BONDI specifications.

Network NUTSTM

for networking, for life

**Fact #1 - Given 570+ Red Hat Certified professionals (RHCE + RHCSS)
In last 12 months, with BEST result in APAC region.**

RHCE | RHCSS | Virtualization

EX-78 | EX-79 | 21.05.09 | EX-83 | 22.05.09

redhat.
PREMIER
PARTNER



North Delhi : 2513, Hudson Lane, Kingsway Camp, North Campus, Delhi-9 Tel. : 65156734, 9310024501

East Delhi : 508, DDA Complex, Laxmi Nagar Distt. Centre, Delhi-92 Tel. : 92448951-2, 9310024502

South Delhi : A-184, Bhaijan Pitamah Marg, Opp. Defence Colony, Delhi-2 Tel. : 46526960-2, 9310024503
www.networknuts.net



You Can
**Touch
The Sky...**



**Let Us
Get Ahead
Together**

Tell Me More About Advertising Opportunities In EFY Magazine

I am aware that EFY magazine presents unspoken advertising opportunities. I too would like to explore the various options available for my organization, so that I can help it grow in terms of sales and brand presence. My details are mentioned below:

Name _____

Organization _____

Designation _____

Address _____

City _____ State _____

Pin Code _____

Phone (OVR) _____ Mobile _____

Fax _____

E-mail _____ Date _____

I would like to have a one-to-one meeting with someone from the EFY team.

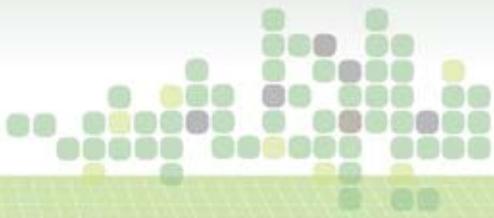
Please send me your marketing kit containing all necessary details.

Please mail/box this form to:



EFY Enterprises Pvt Ltd

D-87/1, Okhla Industrial Area, Phase-I, New Delhi 110020
Tel: 011-26810601/2/3; Fax: 011-26817563, 26812312
Email: efyindia@efyindia.com



Novell introduces service-driven data centre solutions

Novell has announced an integrated vision for the future of the data centre and a portfolio of integrated products designed to help customers increase agility while reducing cost, complexity and risk. The Service-Driven Data Centre provides solutions to build, manage and measure the next generation data centre so IT executives can deliver the business services that end users need, through a flexible, automated and cost-effective infrastructure.

Jeff Jaffe, CTO and executive vice president, Business Units at Novell, said: "The Service-Driven Data Centre delivers an agile, cost-

effective infrastructure through a combination of modular yet integrated solutions for enterprise Linux, virtualisation and workload management, and business service management."

Built on an ITIL blueprint, the Service-Driven Data Centre removes execution risk and manages complexity by establishing a secure, controlled, repeatable and efficient data centre service model that delivers measurable results. Designed to work in a heterogeneous environment, it leverages intelligent management and automation tools to measure performance of a complete IT service against business objectives, explained the company.



Red Hat launches the Teiid data integration project

Red Hat has announced the official launch of the Teiid data virtualisation system project in the JBoss.org community. Teiid is the first open source community project that aims to deliver Enterprise Information Integration (EII) with both relational and XML data virtualisation, according to the company.

"When Red Hat acquired MetaMatrix in April 2007, we committed to releasing the data services technology in the open source community and Teiid is the result of that promise," said Craig Muzilla, vice president of middleware business at Red Hat. "The demand for applications and services leveraging the data stores of a typical organisation never ends. But now, enterprises have a choice between expensive, proprietary data services and an enterprise-class open source platform at a fraction of the cost."

Most open source data integration technologies centre on physically moving or copying data to locations. Teiid bucks this trend by focusing on data virtualisation, which enables real-time access to data across heterogeneous data sources without copying or moving the data from the systems of record. Its Java Database Connectivity (JDBC) and Web Services interfaces are designed to provide straightforward integration with both custom and COTS applications.

Revised Linux certification exams from LPI

The Linux Professional Institute (LPI), a Linux certification organisation, has released new versions of the LPIC-1 and LPIC-2 certification exams. These new exams are available worldwide in English, German and Japanese through the Prometric and VUE testing networks. "This year marks the 10th anniversary of the Linux Professional Institute, and these revised exams are an excellent example of our ongoing commitment to the industry," said Jim Lacey, president and CEO, the Linux Professional Institute.

"During this process, we undertook wide and transparent consultation with enterprise IT professionals and the Open Source community to determine the best placement of exam content within our programme. In addition, we focused on mission-critical technologies, both evolving and established, that will continue to be part of the necessary skillset of a Linux professional." Key exam changes include the following: new content, SQL data management, accessibility, localisation and internationalisation, data encryption, more troubleshooting and security, udev device management, more logical volume management and IMAP/POP, amongst other topics.



Bi3 picks OSS from Jaspersoft, Infobright

Jaspersoft and Infobright, the open source data warehousing companies, have announced that Bi3 Solutions has chosen the companies' joint BI project to build an end-to-end BI suite. The joint, open source project from Jaspersoft and Infobright includes BI, extract-transform-load (ETL) and data warehousing capabilities.

Bi3 Solutions is said to have chosen Jaspersoft and Infobright for their mature product features and flexible integration capabilities, allowing the company to deliver an enterprise-class solution to customers. Jaspersoft and Infobright will provide the reporting, analytics and data warehousing components for Bi3's Virtual Business Intelligence Centre (VBIC), which goes beyond traditional BI applications to provide a fully functional and virtual BI competency centre solution.



Microsoft®

BE LOOKED UP TO

Be the one to save the day for developers with a solution, or share something new to become the one everyone would talk about. Start a new project, join an existing one, or download softwares created by fellow developers.

ENTER A NEW HALL OF FAME AT WWW.CODEPLEX.COM,
THE MICROSOFT OPEN SOURCE PROJECT HOSTING WEBSITE.

CodePlex





KNOW HOW

Q & SECTION

Q A few weeks ago I installed openSUSE11.1, which came bundled with your magazine, on my laptop. I have a partition with WindowsXP, so it installed GRUB with no defects. After a complete installation, I tried to boot Windows from the GRUB menu, but it failed to load and displayed “A data read error occurred, press ‘ctrl-alt-del’ to restart.” Then, after trying a lot, I removed openSUSE and installed Sabayon4; yet still Windows didn’t boot up. I am quite comfortable with Linux, but I need Windows for some games and most importantly, ‘AutoCAD’. Please help me out.

—Adfar Khan, adfar.khan@gmail.com

I would recommend you try re-installing Windows XP again. After reinstalling it, you will not be able to boot into Linux, as your GRUB will be overwritten by the Windows bootloader. Reinstall GRUB using the openSUSE install DVD to configure your dual boot system—check the rescue system options in the DVD.

Q I have started using Mandriva for the last few months. I use it to browse and for other desktop needs. I am very happy with it, as I do not need to update mundane things like anti-virus software, every day. Yesterday I came across the term, ‘shell script’. Please tell me what it means and what it is used for.

—Vishal Abbi, Bokaro

A shell is a command-line interpreter, and a shell script is a script written for the shell. Normally, shells are interactive. It means shells accept commands from you and execute them. You can use commands one by one on the shell prompt. This sequence of commands can be stored in a text file and this text file can be executed by the shell. This is known as shell scripting. There are different shells. All Bourne Shell scripts should begin with the sequence

```
#!/bin/sh
```

On the first line of an interpreter script, following the “#”, is the name of a program that should be used to interpret the contents of the text file.

Q I have installed openSUSE 11.1 along with Windows XP on my old laptop. Everything is working fine except for the wireless. My laptop is a Compaq Presario C300 with a Celeron processor and 1 GB of RAM. Please help me to configure my wireless.

—Pradeep Shakya, Bareilly

I assume that you are able to connect to the Internet using your wired interface. Compaq Presario C300 has the Broadcom Corporation BCM4311 (rev 01) wireless network controller. OpenSUSE 11.1 does not support this wireless card out of the box. To install it, go to <http://software.opensuse.org/search> and search for *broadcom-wl*. You will get the package containing Broadcom’s IEEE 802.11a/b/g/n hybrid Linux device driver for use with Broadcom’s BCM4311. Click on the *1-Click Install* icon and follow the on-screen instructions. Make sure that the wireless card is installed on your laptop by running the following command as the root user:

```
lspci | grep -i broad
```

After this, restart your computer and configure the wireless card.

Q I use Fedora on my desktop. It sometimes happens that when I press the *Eject* button of the DVD/CD drive, nothing happens. I can eject the disk only by rebooting the system. Is there any way by which I can eject the disk without rebooting?

—Ramakrishnan R., Banga

You do not need to reboot your system just to eject a disk from the drive. If pressing the *Eject* button fails, try the following method. Simply right click on the CD/DVD icon on your desktop and select the *Eject/Unmount* option. If even this does not work, then run the following command:

```
$ df
```

Shown below is a sample output for the above command:

```
/dev/hda6           11079832
6877124  3630796  66% /tmp
/tmp/PCBSD7.1-x86-DVD.iso
1902720  1902720      0 100%
/tmp/test
/dev/hdc           693064
693064      0 100% /media/LFYCD_
March08
/dev/hda2           14877092
13689076  420112  98% /home/hdd
```

This tells us that /dev/hdc is our CD/DVD-ROM drive. Now use the following command as the root:

```
# fuser -km /dev/hdc
```

Now you should be able to eject your disk. 

Which one is your career **goal?**



medical



engineering



computer software



computer hardware



hotel management



clinical research

We help you **decide**

Hundreds of Educational Institutions Offering
Professional Courses & Training Under A Single Roof

EduTech's Highlights :

- * An event totally focused on professional education
- * A one-stop facility for students seeking admission into professional courses
- * Counselling by industry top-shots (HR and recruitments)
- * Major plan for advertising and promotion of EduTech Expo 2009
- * Powered by EFY Group's vision to boost professional education
- * Courses catering school pass-outs & fresh graduates

**EDUTECH
EXPO**
2009

13th - 14th JUNE 2009
PRAGATI MAIDAN, NEW DELHI

Laid Off? Why Apply For a New Job When You Can

Be a CEO?



The broader picture: freedom software can turn an engineer into a CEO.

The economy is crumbling. Jobs are being cut. Hope? It's floating. One severely hit sector in times like these is the software development industry. Even companies like Google and Microsoft can't escape the thunder and have had to lay people off in recent times.

Despite this depression, there is a silver lining behind the dark clouds—Freedom Software or Open Source is emerging as the saviour. But it's easier said than done. Let me present a hypothesis here on how Freedom Software saves careers and companies.

The fantastic four

There were four friends working in three different software companies. One fine day, they received pink slips out of the blue. Suddenly, the ones with heavy salaries and the most promising careers were jobless. They met in the evening, concerned about their present and future. They did have some savings and working spouses, so they could survive for a while without jobs. But the economy was not promising anything big.

One of the four friends was into sales and marketing. He confirmed, "Tomorrow, I was supposed to meet a client who is also feeling the heat. The pressure from the US government to cut outsourcing, along with the economic slowdown, was already crushing their outsourcing business."

Apparently the company was waiting for a big contract from their customer, but was required to scale up for that project. To do so, the firm needed to buy new licences for the proprietary products it used. The firm was

reluctant to invest its savings in the current hard times, on something that it was not 100 per cent sure would materialise. Also, the firm could not access funds from banks. If the firm took the risk of investing its savings to scale up, only to find that another firm had out-bid it to win the contract, it'd go bankrupt. Ironically, if it did not scale up, it would still lose the project and go downhill pretty soon.

One of the friends was a software consultant. He offered, "Why don't we meet this client and see if we can help. Either way, we are now jobless and have lots of time to spare."

The four went to meet the client. They surveyed the firm's infrastructure, which was running completely on proprietary technologies. For every new box, the firm would have to buy new licences. After visiting the centre, the four decided to suggest something that could save the firm.

The suggestion

"Migrate to Open Source."

"What?" the owner said. "How can I migrate to Open Source during a running project? And where am I going to get the money for it? And even if it's all possible, who in the world would be able to take up the job of migrating our entire infrastructure?"

The consultant told him that either way he'd have to invest money in buying or renewing licences. If he chose open source, they could bring the cost down dramatically. Most of what he was running had successful freedom software alternatives.

It all boiled down to who would do

the implementation. The four friends had not shut themselves up in the closed-source, air-tight compartment of proprietary software companies they worked for. Instead, they had been using and working on open source technologies all along—thankfully! The consultant suggested, “We will do the implementation, and you can pay us only if your firm bags this project.”

The owner considered this as yet another risk, “But heck, this risk perhaps is worth taking. Besides I’ll still be left with some of the savings.”

So, the four started from the server end and moved towards client side—replacing things piece by piece.

The Windows server was replaced with freedom software servers, and other applications replaced by the open source alternatives. The OpenOffice.org office suite, Apache Web server, MySQL database and GNU/Linux started replacing the proprietary stack.

It took them more than 20 days to migrate the whole infrastructure. On the 21st day, the entire set-up was running on an open source stack, except for a few things that were required by the firm, and that did not need immediate scaling up or were too insignificant.

In the process, the four friends never felt like they were jobless as they were working even harder than before. Now, it entirely depended on whether the firm would get the contract or not, and thus would be in a position to pay them their fees.

Men at work

There was one more challenge for the organisation—it developed custom software and was now using free software to do so. This led to concerns on licensing compatibility issues. Fortunately—look at the irony here—they got to learn about someone recently fired from his job, who was a legal advisor to a software company. They brought him in to advise the firm on the legalities of mixing code.

Post implementation, the friends felt a bit relaxed. However, the next day the owner of the firm called them up asking for their support for one of his clients, who was already using an open source stack and wanted someone to come and fix it—he too was not in a position to afford the fees of his regular consultant.

The four met this client to look into the problem, which was fixable as well. The solution was to replace the proprietary stack with a free alternative. That took a week, but this time they earned some money for their efforts.

On the following weekend, they had nothing much to do. So they planned a vacation for two days. But that was not to be. Their first client apparently had gone around bragging about his successful migration among his peers. So the Fantastic Four got another call at night. An SME was on the verge of shutting down as its licences were expiring and projects were shrinking.

The team of four cancelled their trip and went to meet the folks at this SME. It was an outsourcing firm that provided solutions for a UK-based civic body. There had been a call for using only open source technology in the UK, and the organisation wanted to replace its custom

software with a resource where people could get access to information and file complaints. Now, the Indian firm did not have much exposure to open source, and hence was afraid it would lose a valuable contract.

This was one of the easiest tasks for the Fantastic Four. They already had contacts with a firm that specialised in similar work, and most of the solutions had been released as free software. They scrutinised the SME’s current software stack and found that it could be replaced with the open source stack in a week. They took up the project.

One week passed and they were still working on the project. Meanwhile, they heard from the first client...he had won the project and was now in a position to not only pay their fees, but also make them partners to provide support in the future. They accepted the fee, but due to their involvement in the current project, they asked for more time to think about the partnership.

By the end of this project, the friends realised that one month had passed and they had not sat idle for a minute. They looked back and found they now had good enough exposure to making a living as consultants. They took a decision. The next morning, the four friends officially started an open source consulting firm. They would offer consultancy services on legal, migration, implementation and other matters.

Due to the economic slowdown, more and more firms that were dependent on proprietary technologies were shutting down. The Fantastic Four’s consultancy firm was in demand across the region, as it had already saved three firms from shutting down. But, it was also becoming impossible for them to be at the clients’ sites all the time. They hired two engineers who had recently been sacked. Then they hired one sales person to take account of all the clients. They also needed another database migration expert.

Feeling stressed out?

It was just three months since the four had got ‘pink-slipped’. They were going to need more people, given the rate at which projects were coming in. And they sat and reflected—three engineers and a sales person who had been fired and had been so unsure about the future, were now company owners. They had saved three companies from shutting down and had even hired eight engineers, while the IT giants were busy laying off staff.

It's only a hypothesis

Well, I started by saying this was just a hypothesis. But hey, there is a possibility that it’s already happening somewhere—you never know. 

By: Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company’s ‘paper armours’. He is a big movie buff and prefers listening to music at such loud volumes that he’s gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.



Is It a Bird? Is It a Plane? No! It's PC-BSD!

The FreeBSD projects boast of the most stable operating system kernel in the world. Created from the University of Berkeley's BSD4.4Lite sources, it's a valid claim. And when such a kernel is blended with one of the most comfortable userlands of the world, magic happens. Ladies and gentlemen, we bring to you PC-BSD, a perfect fusion of BSD's kernel and GNU's userland—an operating system that's aesthetic, pleasurable and complete in every aspect, and which has been designed keeping the assassination of Microsoft in mind.

*J*ust about a week back, I asked my father, "Why don't you use Linux?" He replied with another question: "Why do people use Windows?"

As much as we may rave about it, laymen are just not enamoured by or convinced about Linux. We know that the days when Linux meant staring at a shell prompt and typing out cryptic commands, are long gone. Nowadays, most Ubuntu users don't know that they are using GNU/Linux. They find it much more aesthetic, usable and logical than Windows. They have no idea that under the

skin of such an operating system is a heart so powerful and complex that it takes nothing less than a power-cut to bring the system down. No more black screens, no prompts, an integrated and tightly knit GUI... ladies and gentlemen, the future of UNIX computing is here.

Let's just forget Linux for a bit and, instead, face the ultimate truth. We are using a UNIX-like operating system. Yes, that's right. UNIX was born in AT&T's Bell Laboratories as a system written in Assembler for an ageing PDP-7 system. After



Figure 1: Default desktop

the complex OS got too difficult to maintain in Assembler, the C language was created and UNIX was re-written in C. By about 1977, Berkeley Software Distribution (that's BSD) started as a set of add-on software to Sixth Edition UNIX. Ultimately, by the time BSD4.4 came out, it had forked off from AT&T's UNIX and was on its own. However, 4.4 was the last BSD version, because BSD had run into licensing and patent issues. After the legal battle, 4.4 came out in 2 versions—4.4 Encumbered, with a lot of AT&T and some proprietary code, and 4.4Lite, with only the Berkeley code.

4.4 Lite had huge chunks of the kernel missing—critical chunks, which made the kernel unusable. Two major projects started off, and each took the 4.4Lite code and started developing on their own. We know of them as NetBSD and FreeBSD. The former specialises in security, while the latter boasts of stability. Around the same time, Andrew Tannenbaum developed MINIX, a 16-bit UNIX clone to teach students operating system design. MINIX had nothing to do with UNIX sources. And, watching Andrew, Linus Torvalds came up with a 32-bit clone, called Freax. Freax did not have anything to do with either UNIX or MINIX sources. And after Freax was complete, the manager of Linus's FTP site changed the name to Linux, because he didn't like Freax. Thus we had five UNIX and UNIX-like operating system kernels: UNIX itself, NetBSD, FreeBSD, MINIX and Linux.

Bang Bang Bang—PC-BSD is here

Mac users love their OS, and Mac OS X is built on part of the FreeBSD 5 kernel and that entire userland. So some guy called Kris Moore took apart FreeBSD 6.0 (at that time it was the latest) and built his own OS around it. Version 1.0 was released on April 26, 2006. The OS aimed to be the most Mac-ish, using a BSD kernel, a highly customised version of KDE 3, and a revolutionary... ahem, Mac-like package management system.

Cut to the present: I'm sitting down in front of Mozilla Firefox 3.0.8, running on PC-BSD 7.1 Galileo Edition, based on the FreeBSD 7.2 pre-release kernel,



Figure 2: My desktop after changing a few things

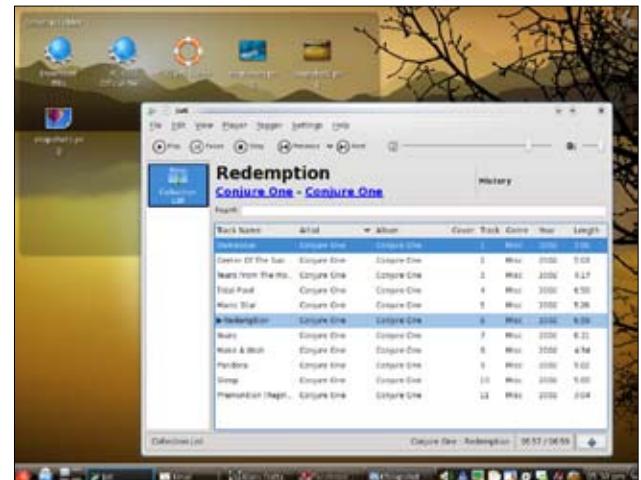


Figure 3: Enjoying my music on JuK

and running KDE 4.2.2. Savvy?

Yeah! PC-BSD is refreshing after the bores and chores of Linux. It comes on a plastic platter 15 centimetres in diameter, and to get it started, you install it. To do that, you must pluck out some sort of a tray from the front of your PC, put the platter silver side down on that tray, and close the tray. Then restart your computer, and follow the steps on that huge box that contains installation instructions.

After the actual installation is complete, the first boot needs some attention. After X starts up, a pop-up comes along, telling me to supply some information about my graphics card. I have an NVIDIA GeForce 7100/nForce 630i, and I don't like what I see. The string "vesa" is staring at me when I look at the field specifying the video driver. Not good. So I open up the drop down menu, and see three NVIDIA drivers—version 69, for legacy cards; version 173, the stable drivers; and version 180, the current drivers. I select version 180, and X re-starts. WhoHoo! I'm finally staring at KDE 4.2.2. At this moment on, there really is no difference between PC-BSD and Kubuntu. Well, save one: you don't have Apt here.

The software set is pretty standard, bundling in the

Install experience

Installing PC-BSD from that round platter is a no-brainer. But, of course, you can't really live without a brain. But then again...

1. Here we go. After you insert that platter and re-boot (no, not throw a boot, throw them at politicians) your PC, a text screen shows up. Scary as it may seem, hit 8 and press *Enter*. (Or just press *Enter*. You cannot use ZFS then.) Hit *n* when it asks you if you want to check the integrity of the installer images.
2. Up comes the first installer screen! Leave the keymap and language settings alone, but change the timezone to Asia/Calcutta (that's the only IST zone present). Hit *Next*.
3. Agree to the licence if you want to go ahead. There's a lot of proprietary cruft in there, though.
4. In this screen, leave all the settings as they are: New Install, Desktop Version, and CD/DVD source. Hit *Next*.
5. Here, fill out your user profile details. After you're done with the text entry, change the shell to */PCBSD/local/bin/bash*. That's the only way to survive. Hit *Add*, check the auto-login details and hit *Next*.
6. Here you get to choose your hard disk details. Choose a hard disk in the first selection box, then a partition on the second, or just select "Use Entire Disk". If you have less than 1GB of RAM, ZFS won't work (it just won't show up). Here, UFS+SU is the best choice. Review your choices, and hit *Next*.
7. In this screen, you get to choose optional components. At a bare minimum, you need Firefox, because Konqueror (included as default) is not good enough. Choose Amarok if you must, or stick with the included JuK. Hit *Next*.
8. After a bogus *Ready To Install* screen where you just hit *Next*, everything is installed.
9. After the first boot, a small applet comes up asking for videocard details. Choose as per your own hardware, and then, you're done! As simple as 1-2-3! Um... add 5,6,7,8 and 9 there ;-)

entire KDE 4.2.2 application set. At installation time, you will notice that a lot of stuff like OpenOffice.org, Firefox and Amarok were optional. I did install all of them. Amarok and OOo versions come up as shockers: 1.4 and 2.4, respectively. Everything else is fine. And the moment it booted up (it was five days from the release date), an update for each and every optional package was available. A 500MB download, all in all. OOo 3.0.1, Amarok 2.2, what else? Those guys could have populated the DVD with those instead of legacy versions.

Take it for a ride

I did that for starters. There were no disappointments, save some hiccups. The first thing I did was remove the normal launcher and enable Lancelot. An annoying problem was that while log-out works, hovering the mouse on 'Switch

Users' brings me a configure GDM menu item.

I next tested out the bundled version of Wine 1. I browsed away with Dolphin and double-clicked on the installation file of Foobar2000, a free (not *libre*) music player for Windows. It installed fine, but would just not play any audio. WinCfg didn't help, and running Wine from Konsole resulted in a segfault. Bye, bye, Wine! By the way, it's supposed to run Linux binaries natively, but I didn't check.

Broadband connections in Bridged Mode (that is, the dial-up-like mode) would not work, so I ended up hacking my modem and setting it to 'always-on' mode. Now the Internet rocks! It's fast, and it's quite stable.

Speaking of speed and stability, PC-BSD is by far the most responsive system ever to run on my PC. It's stable as well. Save for the Wine segfaults, nothing has gone wrong yet. 3D effects are enabled by default and they rock. MP3s play by default, as do H.264 video. By the way, PC-BSD is not for FOSS purists, because starting with the BSD licence itself (which makes it 'almost public domain', save reserving a copyright), it includes all sorts of encumbered drivers such as NVIDIA's, and codecs such as MP3 and H.264. By the way, Flash is included as well.

The nitty gritty

PC-BSD wouldn't install on logical drives, but the BSD people say that it should install on an extended partition, preferring their own slices and disk labels rather than Windows's LDs. Coming to partitioning, I must share my installation choices with you. I booted up the installer with option 8, or with support for Sun's Zetabyte File System (ZFS). I selected this in the partitioning screen, and the installer barked at me: "ZFS is for 64-bit systems with more than 4 GB of RAM." I have two. Who cares about warnings? I ignored it and, ooh la la! ZFS is literally a rocket at managing the storage of my files!

PC-BSD is geared towards newbies, but when they see the black screen with three lines of text after starting their PC (the boot-loader), they may faint. If they don't, after pressing the desired F-key for FreeBSD (yes, it's not PC-BSD), they'll come to a screen with pretty nice choices that are also baffling. So what should they do? Just hit *Enter*, guys—the default choice is fine. After that, don't try to make any sense of those meaningless lines of text flashing by. And don't enter your login name when prompted; leave it as it is. KDE will start automatically.

Just how easy is it to crash this stuff?

It's actually very easy to crash it, but it takes a while to figure that out. Actually, I did it accidentally. I was just casually browsing the directory tree, and chanced upon a boot-splash image. I knew there must be a way of enabling it. Turned out there was: four commented lines in */boot/loader.conf*. I uncommented that stuff and rebooted the PC. That was the last time I saw the desktop.

Of course, you could use the classic *rm -rf /*. The Avada Kedavra command!

Back to the stable: The Push Button Installer

That's our package manager. It's called PBI, and it stands for Push-Button Installer. There's no automated download-and-install system here, it's all a-la Mac. In Macs, you download a DMG file. Here, you download a PBI file. The PBI file is self-contained, has all the dependencies packed in, and installs to its own prefix, under the */Programs* directory. Double click the PBI file, and after entering the root password, the first of the installer screens appear.

Here I needed to install *Last.fm*'s Qt4 radio app. To start with, I headed to www.pbidir.com. Then I clicked on *Multimedia*, and then *Last.fm*. Select *32bit* on the first page, then a mirror from the list (I use *isc.org*) on the next, and then finally click on the download link on the third page. Brew some coffee while this 14MB file is being downloaded.

After downloading is complete, double-click the file. It's nothing but 'click-next'. You'll be done in a jiffy. Menu entries will be created in Lancelot and the launcher, and an icon will appear on the desktop. We are done! Way too simple. Of course, you might need to get something other than *Last.fm*. PBI dir is the official PBI repo, and you'll get everything from it. The version of PC-BSD included with *LFY* is 32-bit. So always get the 32-bit version.

Well, what about that cup of coffee? Have it while you channel some instrumental music from *Last.fm* over your 7.1 channel speaker system. You deserve it.

Last thoughts

You guys know that I suffer from Distromania, and I have already left PC-BSD in search of a better operating system. But to me, PC-BSD left a deep impression for being a stable, logical and easy-to-use

operating system based on a rock solid kernel that's impossible to crash (unless you tweak your *loader.conf* or run the Avada Kedavra command). It's fast, responsive, and has a host of software available for it. Everything works out of the box, with minimal configuration, and no googling or trouble-shooting.

NVIDIA drivers are available out of the box, MP3s play, as do FLV movies. Everything looks sleek, and works together. I'll let out a little secret. I'm actually on assignment here. I was told to review Zenwalk, but I just couldn't even boot it. X refused to start with vesa drivers, broadband went haywire, LILO ruined my bootsector and basically made my life hell for the three days I struggled with it. At the end of it, I put my hands up and surrendered, and they gave me PC-BSD. PC-BSD is the exact opposite of Zenwalk. Everything works, and works harmoniously. After you install this OS, you'll love your brand new PC!



PC-BSD 7.1



Pros:

Speed, stability, ease of use, package management, proprietary hardware and software works out of the box.

Cons:

Because it's based on BSD, Linux command-line fans will have a lot of trouble. Plus, OOo 3 and Amarok 2 are not included, but are available as upgrades.

Platform:

x86

Price:

Free (as in beer)

Website:

www.pcbsd.org

By: Boudhayan Gupta

Boudhayan is a 14-year old student who suffers from an acute psychological disorder called Distromania. He owes his life to Larry Page and Sergei Brin. Apart from that, he enjoys both reading and writing, and when he is not playing with his Python ;), during most of his spare time he can be found listening to Fort Minor, or cooking.



TRAINING PARTNER

RHCE / RHCSS Exam Centre

Hundreds of hardcore Industry Professionals Vouch for it.



RHCE only @ ADVANTAGE PRO

- Pioneer in LINUX training since 1999.
- Chennai's #1 RHCE training Centre
- 100% results every month.
- Many centurm students.
- Also get trained on MySQL - CMDBA, CMDEV, PHP, Perl, Ruby, Python, Ajax...

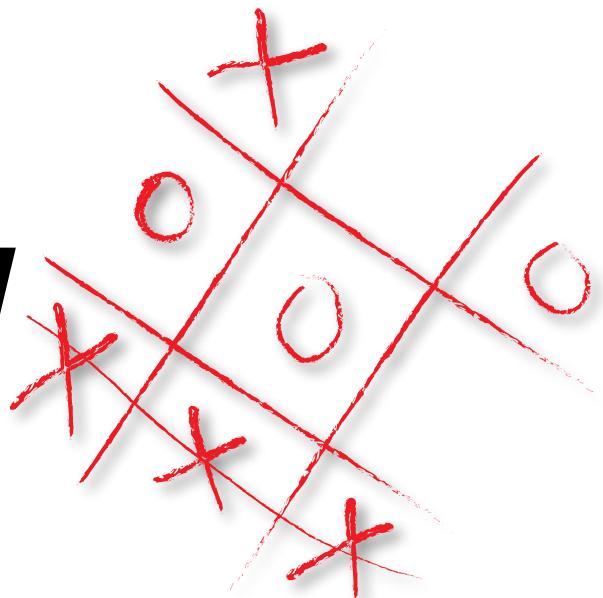
**Next RHCE/RHCSS Exam on
11th & 25th May 2009**

**Don't wait.
Register Now.**
Let the success vouch for you!

VECTRA TECHNOFT **ADVANTAGE PRO**
IT TRAINING DIVISION OF
 Vectra Technosoft Pvt. Ltd.
Open Source Academy

**Regd. Off: Wing 1 & 2, IV Floor,
Jhaver Plaza, 1A, N.H. Road,
Nungambakkam, Chennai - 34.
Ph : 28263529 / 3530 / 3540
Telefax : 28263527
E-mail : enquiry@vectratech.in
www.vectratech.in**

Still Radically Simple



...but is that enough? PCLinuxOS 2009.1 has finally been released after a wait of almost two years. We take a look at how the new version of this former Distrowatch topper fares.

T’s only fair to first mention a few words about myself. I have been a hardcore Windows user, until about two years back. Since then, although I have been running Mandriva on my machine, I was still predominantly a Windows user because I am a professional gamer. As typical with most OS migrants (especially from Windows), KDE was my very first choice to test drive, and I have stayed glued to it till date. Hence, it’s no wonder that I was asked to review PCLinuxOS 2009 the moment our IT admin downloaded it (owing to its similarity with Mandriva, and KDE runs on it, by default).

Background check

PCLinuxOS is a GNU/Linux distribution that was built on Mandriva and was launched as a set of RPM packages. Bill Reynolds (a.k.a. Texstar) had created PCLOS as a fulfillment of his wish to package source code without having to deal with the rest of the world. It later evolved into a complete desktop operating system with its own unique set of features. The distro’s stability and ease of use made it popular pretty fast and even topped the distro ranking at Distrowatch in 2007.

Configuration of the test machine

I ran it on a pretty low-end laptop, a Compaq Presario C300, with the following specs:

- Motherboard: Intel 915
- Processor: Celeron M 1.67 MHz
- Memory: 512 MB DDR
- Hard Disk: 60 GB SATA

The installation

After I booted the machine off the LiveCD, I was presented with the option of a LiveCD install and a few safe modes. There was no option to install PCLOS 2009.1 right away. Anyway, before loading the operating system, I was asked for my keyboard layout and was directly presented with the log-in screen.

The default IDs are as usual, ‘guest’ and ‘root’, with their respective passwords mentioned at the top left of the screen. I logged in as the root user right away because I wanted to install this baby on my machine, ASAP.

While you log in, you are presented with a loading screen that actually shows various icons symbolising each of PCLOS’ major releases. Although I am not too confident if it is a welcome change from KDE’s traditional loading screen, I am sure I would

prefer the latter, any day. What's more, there's bad news for all you KDE4 enthusiasts. PCLinuxOS 2009.1 does not ship the latest version of KDE. PCLOS runs on KDE 3.5.10 since the creators still do not have enough faith in the latest version. However, they have promised to make it available on the repos, very soon.

When I clicked on the installation icon on the LiveCD desktop, it was interesting to find a message window asking me to remove a few video drivers from the installation if I do not need them. For example, if you have an ATI card, NVIDIA drivers would be redundant and it's better not to install them in the first place. Removal of unnecessary drivers also gives you a faster booting sequence. Of course, if you are unsure, you can always click on 'Cancel' and move on with the installation process.

PCLOS has a pretty simple drake-live installer that doesn't ask for many options and gets the job of installation done in about 10 minutes flat. However, the partitioning window is not too newbie friendly, especially considering that PCLinuxOS is targeted at people who wish to migrate from other operating systems. Moreover, PCLOS gave me the option of creating a custom partition or installing it on available free space, but there was no option to install it on the complete disk. This just makes one wonder, are the guys at PCLOS trying to tell me that it is always good to have another distro installed in my computer, as a back-up? Don't they have faith in their own offering?

My computer already had openSUSE installed in it and I wanted to use the same 'home' folder. It was great to notice that PCLOS did not offer to format 'home' by default, but only the 'root' drive. After partitioning, the operating system was installed onto my machine and I was presented with the GRUB bootloader settings. While you can always tweak it out, for the benefit of the newbies, you can just keep clicking *Next* till the installer asks you to remove the disk medium and reboot the system.

Oh, I almost forgot, there is also a PCLOS-GNOME version available for GNOME enthusiasts.

Initial impressions

One of the first setbacks I experienced was when I found out that there is no 64-bit edition for PCLinuxOS 2009.1. What's more, booting into the operating system was a rude shock because of the really ugly artwork that they have come up with. Although a customised PCLOS button is always a welcome change from the same old KDE menu button, I only wish they had used a bit more aesthetics while designing it. Moreover, the default theme used by PCLOS will actually make the windows look like that of Windows Vista. Nice try, I'll say. The very first thing I did, after logging in, is change the theme to 'Plastik' to give myself some visual respite.

What also irritates me, at times, is why the distros don't manage to automatically gauge a machine's graphics capability and switch on the 3D graphics automatically? For new users, there is a high probability of them using the same distro for years, without even being aware of its 3D capability,

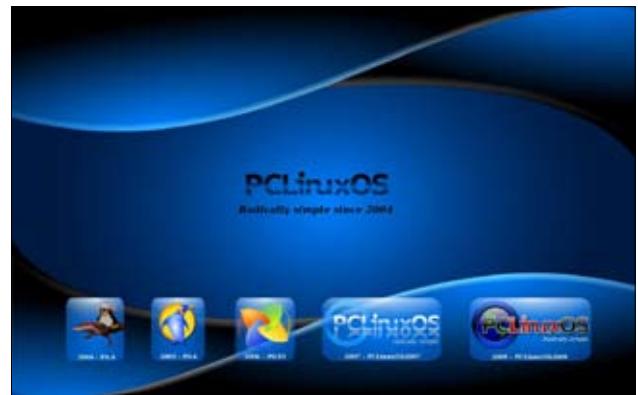


Figure 1: Revamped KDE splash screen



Figure 2: Default desktop

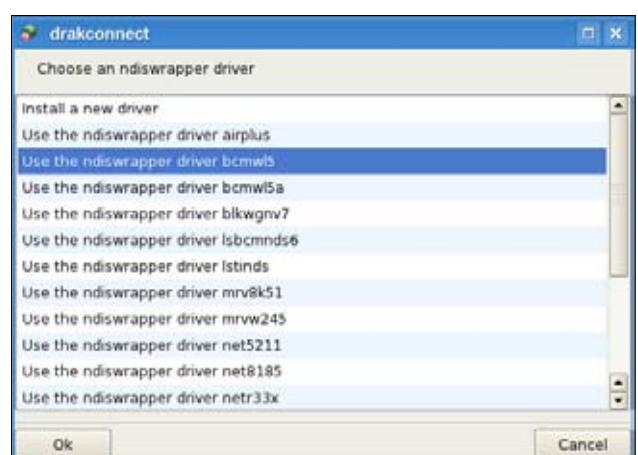


Figure 3: Configuring wireless

ever! At least, a user should be presented with the option of enabling 3D graphics right after logging on, for the very first time. Well, speaking of PCLOS, activating Compiz in my machine was a farce since it degraded the performance to a frustrating level. But, then again, I guess this time the fault is at my laptop's end—it's too old.

A deeper look into the distribution

The software manager in PCLOS is pretty confusing, at least for me. Allow me to present my case. Firstly, Synaptic as a package manager, for a RPM distro, hits you right at the middle of the head. What's more, one can't even install an



Figure 4: An ideal distribution for internet PCs

RPM file by double-clicking on it. I had to download KPackage to be able to install RPMs without any hassle. (*Editor's Note*: Boy, and I thought Synaptic as the default package manager was a seller—at least for most people ;-))

The repositories have really old software. It beats me that PCLinuxOS still uses printer-drake, something that has been discontinued even by Mandriva, which maintained it. Unfortunately, I did not have a printer nearby to test its functionality.

Fortunately, there's something good about the distro, too. PCLOS Control Centre shows some good improvements, including changes to how networking, firewalls, proxy and shares are handled. Software load times are very good, at least compared to Mandriva and openSUSE. Moreover, since I use the 'printscreen' a lot everyday, finding the button automatically configured with KSnapshot brought a smile in my face, right away.

Bundled software

The desktop version of PCLinuxOS 2009.1 has exactly the same number of software as on the LiveCD. Moreover, it is just the sweet spot for the 'migrators' since the distro doesn't pack in as many software as in Mint and Sabayon, but neither is it as bare as in MiniMe.

One function that won my heart

is the option to create a LiveCD and LiveUSB on-the-fly. While I created a LiveUSB of PCLOS itself, I have a feeling that the tool shall allow one to create a LiveUSB of any distro, since it asked me for the image file of the operating system.

Compatibility issues

I, personally, didn't find a single compatibility issue with PCLOS. Every single piece of hardware on my laptop got detected seamlessly by the operating system. Besides, since I need to use ndiswrapper for the wireless adaptor, PCLOS actually gave me options to choose the one that's for me and went on to install it without a hitch!

I just wish the DVD drive of my personal computer was in working condition, because I was very interested to check out how easy/difficult it gets to detect my Web cam.

Multimedia capabilities

I must admit, I was blown away by the distro's multimedia capabilities. It didn't matter which audio or video file I threw at it, PCLOS managed to play it with absolute *elan*. All my MP4 and AVI files worked out of the box. It was also good to see KMplayer as the default video player, and not Kaffeine. My only complaint... I wish VLC and SMplayer were also bundled with the distribution.

PCLOS for an Internet PC

PCLinuxOS 2009.1 is the perfect GNU/Linux distribution to be installed in an Internet PC. With the latest craze for this genre of computers and with the attempts to reduce their retail prices, PCLOS should be an automatic choice. It comes with the latest versions of Mozilla Firefox and Thunderbird. That takes care of the browsing and e-mail needs. For IM, there's always Kopete to bail you out. Yet, a desire for Pidgin does keep simmering. What's more, it also has Flash already installed in it. So, one can start YouTubing right away, without the hassle of downloading the plug-in separately, as is the case with almost every other operating system! Also, apart from KTorrent,

PCLOS also has FrostWire to take care of your downloading needs. However, beware! Just like proprietary software, unauthorised downloading of movies and audio files is still a criminal offence.

For a distro version that took two years to launch, the level of development is zilch. Moreover, I found PCLinuxOS 2009.1 a bit behind the times in a few areas.

Rounding off

Although PCLinuxOS 2009.1 is absolutely solid and recommendable to one and all, the latest artwork is an absolute turn off. I have been using Mandriva 2009 for quite some time now, with KDE 4.2 on it, and had wanted to shift to another distro for a few weeks. However, the very looks of PCLOS 2009 forced me to stay put with Mandy and bide my time for the Spring Edition. Otherwise, if you are looking out for just plain functionality and do not give a hoot about how your operating system looks, you should be one very happy PCLinuxOS user, for sure!



**PCLinuxOS
2009.1**



Pros: Painless Installation, pre-loaded ndiswrapper drivers, well-organised Control Center, Software bundle, Audio/Video Codecs Support, Rock Stable.

Cons: Pathetic Artwork, RPMs can't be installed with double-clicks, unoptimised Compiz, KDE 4.2 absent.

Platform: x86

Price: Free (as in beer)

Website: www.pclinuxos.com

By: Sayantan Pal

An avid Twitter user and a social media enthusiast, the author is a passionate blogger and a professional gamer too. He also feels compelled to be opinionated about anything that comes his way, be it Linux distributions, our marketing strategies, table etiquettes or even the fabled Ramsay movies!

We have crossed **2.5 million** hits per month !



Join These Growing Numbers And Benefit...

electronicsforu.com A comprehensive portal on electronics

Popular sections



Circuit Lab

Hundreds of circuit diagrams of different projects, for **FREE**



Top 100

Most popular technology websites, measured by their hits



Events

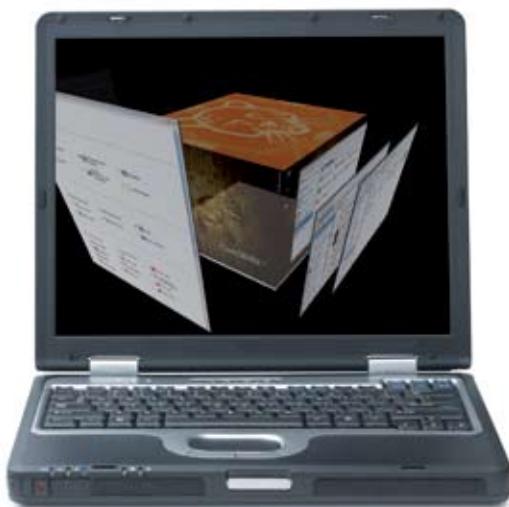
A comprehensive list of technical events, in India and abroad

Looking for something related to electronics? Try electronicsforu.com. This website has circuit diagrams, technical articles, bulletin board, dedicated index of electronics websites, list of popular websites ranked by hits...you name it, it's got it.

Check it out today!



EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020
Phone: (011) 26810601-03 Fax: (011) 26817583/26812312 Email: info@efyindia.com



Kaptan At Your Service

What do you expect from a perfect desktop operating system? Things to work out-of-the-box, I suppose. Well, Pardus 2008.2 works like a charm!

We, the distro hunters, face a drought season twice a year—during autumn and spring—when the development teams of the four major distros that most of us care about are busy fixing the last-minute bugs to meet their release deadlines. As you're reading this, Ubuntu and Mandriva are already up for grabs, but the final versions came in too late for us to include them on the optical discs. So, we decided to make a compilation like we did back in October 2008 with mini, non-so-well-known distros to fill in the media.

The plan was similar, but there's a slight twist this time. There's a PC-BSD ISO image in the DVD, and there's also a PCLinuxOS live-cum-installer OS for a CD. Thankfully, PC-BSD doesn't take up the whole of the DVD, which left us with some space to fill, and we finalised on including ISOs of three mini distros—SimplyMepis, Dreamlinux

and Pardus. After assigning others the responsibility of reviewing four of the distros, we were still left with Pardus. Oh well, I got stuck with it as you can see. *So, I said to myself, I said... damn!*

I'm a Mandriva fan at heart, and have been flamed for trashing every other distro to the point that I'm always the last person the team would give the responsibility of reviewing a distro. Good that most of my comments about a distro are rendered not suitable for print. Sorry folks! Everyone has a bias.

Enough already! Would you mind doing your job? Yeah, yeah...!

What on earth is a YALI?

Yet Another Linux Installer, as they like to call it, is the default installer for Pardus. They say it's written in Python and Qt—*like it makes any difference* was my first response. Turns out it does—or maybe it's just the UI team



Figure 1: YALI's partitioner in action—notice the help window on the right

that matters here—because YALI is seriously quite good, someone opined.

It's the usual chore of triggering your BIOS to boot from the CD/DVD ROM, placing the CD (yeah, don't forget to burn the Pardus ISO image available in the DVD to an empty disk; I just noticed that the other distro reviewers this month totally missed this point of providing the complete picture about stuff that matters ;)) in the drive, and hit the reset button of the computer tower. (*Dude, it's called a cabinet. Oh, shut it!*)

And thus, the journey began! First looks are impressive—the colour scheme of a yellow-orange shade, I've got to say, is very good. *Wait a sec now.... Good gawd! What's with this foreign language on the boot screen?* Turns out it's Turkish, by default. Thankfully, there's an F2 button on my keyboard and I could switch to English. *Yes, that's much better now!*

Then press the F3 button to select the video mode as well. Since my CRT monitor supports up to 1280 x 1024px, I decided to give that a try. The default option to boot Pardus seemed good enough to me, so I hit *Enter*. The nice and plain splash screen with a progress bar that basically hides the verbose kernel boot messages was soon replaced by the installer's welcome screen with the licence message. It's GPL, *phew!* Obviously, they needn't ask me if I'm okay with it or not, but good that they have put a button there if anyone wants to read the GPL—it's good information for some folks. The addition of a *Release Notes* button is a bonus, and I'd say has been implemented at the right place. Not to forget the *Show Help* button as well—upon clicking it, a nice help box on the right shows up with information on what to expect (Figure 1). Nice add-on for the ones in need of help, as they can leave this on throughout ;)) Anyway, overall, it's a very well-thought layout for a welcome screen.

Next up is the CD integrity check screen. *Useless!* Then the keyboard layout—English is good. Then comes the time zone—no world map here, so scroll down in search of 'Asia/Kolkata'. Then the screen to create the system

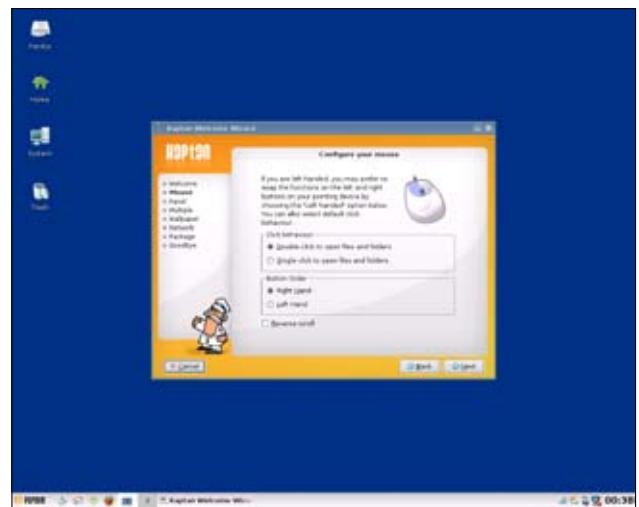


Figure 2: The Kaptan guides you after first boot

(I mean ordinary) users—you can create multiple users, give them admin privileges, and set it to auto login, if you like. Set the root password and host name in the next one, following which the partitioner pops up.

OK, it lists both my hard drives, and there's no way I want to go with an automatic partitioning scheme, so I chose the time-tested manual option. The following screen lists both the hard drives as tabs, and presents a horizontal bar as the partition (Figure 1). This is where you'll select the chunk which says 'free space' (or a pre-partitioned space that you can format), move the slider to set the size, click the drop-down menu to set the 'use' criteria—root, home, swap, etc—and hit *Apply*. The labels for 'use' cases are interestingly titled 'Pardus System Files' for the root FS, 'User Files' for the home partition, 'Swap Space', and 'Storage Area' for God knows what (if you happen to figure it out, do let me know because I'm too lazy to test it). I figure these labels would make more sense to new GNU/Linux users, although they might confuse a veteran. Nonetheless, the partitioner seems very intuitive, and looks much more useful compared to the partitioners available even with some of the major distros.

Next up are the bootloader options, followed by a final summary screen listing the details of all the selections you've made in the installer. Interestingly, YALI is yet to make any changes to your system; so this is your final chance to go back and make amends if you suddenly notice you've done something stupid during the previous steps. If you're satisfied, you can hit—well, *Next* is now greyed out, so—*Begin Install*.

Following this, YALI will partition/format the disk, and proceed with the installation of packages. Yes, you're right—there's no package selection option; everything will be installed by default. The installer is not the fastest I've seen—took around 20 minutes on my Athlon X2 5800+ system with 2gigs of RAM, and around 30 on my Celeron laptop with 512 megs of RAM. But the information on the installer presents a case for lots of software, so if you're

new to GNU/Linux, this is your chance to learn. Finally, it's "Goodbye from YALI" and you've to hit the *Next* button to reboot the machine to "Enjoy your fresh Pardus!"

Bottom line: Although you only install an OS once, YALI still leaves a good impression. And I hate the fact that I won't be using it as often as I would use, say Firefox or Amarok. However, I would love to use it over and over again on other people's systems—*anyone game for it?*

Kaptan helps you get started

The GRUB boot screen after install will list all the other distros installed on your system. The system boots to the KDM screen (if you haven't set the user on auto log-in) seriously fast, thanks to MUDUR, the Pardus home-grown *init* system. The KDM screen looks customised with the only options being: enter user name, password, login or quit to shut down/reboot/suspend the system. This is where I realised that hitting F3 and setting the screen resolution earlier had set 1280x1024px as the default resolution post installation—quite impressive, considering it's using Vesa, and not the proprietary drivers from Nvidia.

Upon your first login, the Kaptan greets you. Kaptan is nothing but a *click-next* wizard to set up your desktop. It starts off with mouse settings (Figure 2), then the desktop layout style (placement of panel, and the ability to replace the default KDE3 menu with Kickoff), followed by setting the number of workspaces. Then comes the wallpaper (by default there is no wallpaper—I chose Canis aureus 02, the one for this release, which has a nice artwork of a running wolf, fox, jackal, or something similar.). Then came network set-up.

This is interesting, as on my laptop, it listed the wlan0 interface, and allowed me to configure it out of the box, with no requirement to load the Broadcom drivers first. The final option before a 'Goodbye' from the Kaptan is the ability to configure package manager settings. This is where you can set the parameters for updates, like show an icon on the system tray, the interval to check for updates, and add the 'contrib' repository. Every time it requires a system-level change like while configuring the network or setting 'contrib' repository in the package manager, it prompts for authentication.

In fact, the final 'Goodbye' screen is handy too. It has the option to reconfigure your graphics. While testing it on VirtualBox, it even reported that the resolution was too low (being at 800 x 600). I noticed a second spelling mistake here—the first one was somewhere in the installer. It said, "You can change your resolution *bye* using Display Manager." On my AMD system, launching the display manager prompted me that I might be better off using the Nvidia v177.x driver in my system, instead of the default Vesa. Later on I found that v180.x is also available in the software repository, and although I had installed it later using the package manager, the display manager continued to report that I might be better off with v177.x. Strange.

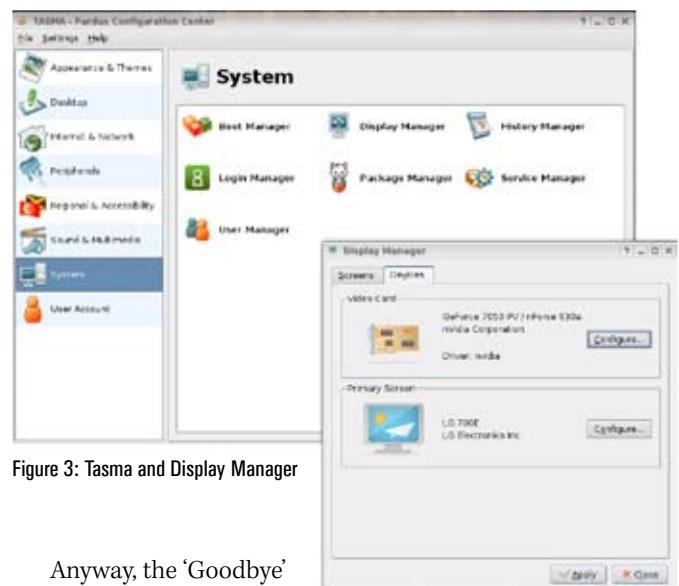


Figure 3: Tasma and Display Manager

Anyway, the 'Goodbye' screen also introduces you to Tasma, the Pardus configuration centre, and the *User Group* and *Help* pages. Upon clicking the *Start* button of the latter it launches Firefox to take you to the 'Contact' page of the Pardus website, which lists information on mailing lists, IRC, Bugzilla, etc.

Tasma helps you configure

At first glance, it looks like the Pardus team has re-branded KControl to Tasma. However, unlike KControl, the home screen explains the purpose of this tool. The broad categories on the left panel are the same like KControl, and so are the modules for these individual sections. Oh, wait a second, a lot of third-party modules have also been integrated into it, and there're more modules to be accessed using *Settings*→*Show Extra Modules*. These essentially make it similar to the Mandriva Control Centre or SuSE's YAST.

One very cool module I chanced upon here is the *History Manager*. It lets you take snapshots of the system and roll back to a previous system state if something goes wrong. Although I didn't get a chance to break the system and then try this out, my understanding of this tool is that, besides giving you the ability to take snapshots of the system state manually, it also keeps a log of package installs and removals as per their date and time, and thus essentially gives you the privileges of rolling back to a previous state.

Other modules include: user configuration, firewall setup, network manager, desktop sharing, display manager, TV card configuration, service manager, package manager, etc. Refer to Figure 3.

Apps helps you get the work done

Back to the desktop, I like the fact that the developers haven't overloaded the desktop with too many shortcut icons. In fact, by default, it's just the *System* icon (which is configured to load the sysinfo KIO slave), *Home* and *Trash*.

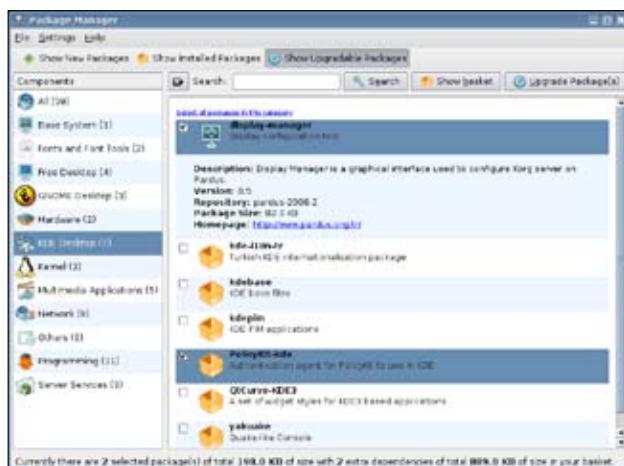


Figure 4: The Package Manager has more to offer you

Of course, it'll also display the icons for removable devices when plugged in.

While we are on the subject of icons, Pardus uses something called the 'Tulliana' icon set by default, instead of the hideous-looking Crystal that comes with KDE3. I find Tulliana quite gorgeous.

The Pardus menu is a customised KDE menu, which only lists five options under the 'All Applications' category—*Tasma*, *Package Manager*, *Home*, *Find*, and the *Programs* menu. This *Programs* menu, in turn, lists the software section inside it, like *Games*, *Graphics*, *Internet*, etc, which otherwise show under 'All Applications' in KDE3, by default. I don't know how usable this customisation is, but it probably makes the lives of those used to the Winduhs ways of doing things, a bit easier.

When it comes to applications, Pardus has everything covered for a typical desktop user—from multimedia and graphics apps to Internet and office tools. OpenOffice.org is still at v2.4.1—hope the next Pardus version carries v3.0 along with KDE4.

Anyway, this distro even includes all the multimedia apps like MP3, Flash, DivX, encrypted DVD decoder (*libdvdcss*), etc, by default. Makes the life of a new user easier, wouldn't you agree?

Something strange... in OpenOffice.org the English dictionaries are missing—so there's no spell check. I tried to look it up on the package manager, and had to install the *openoffice-dicts* package to get it working. This should have been installed by default.

It's also odd that the MPlayer UI is not translated to English. My cure is SMPlayer, because I anyway don't use the MPlayer UI.

Obviously, you can install more apps using the Pardus package manager—*what's it called again?*

PISI helps you get more stuff

PISI—Packages Installed Successfully, as Intended—is what they like to call the package manager. The command-line syntax is very similar to apt. So in order to install a package (e.g. *kdesdk*), you need to do the following:

pisi install kdesdk

To search for a package—e.g., Firefox:

pisi search firefox

To remove a package—e.g., Firefox:

pisi remove firefox

To update the repository database:

pisi update-repo

There are many more options, so take a look at:

pisi help

Of course, there's also a GUI front-end. You can access it either from the Pardus menu, or the icon sitting next to the clock. The package manager has an innovative UI. There're three tabs on top: *Show New Packages*, *Show Installed Packages*, and *Show Upgradable Packages*. Take a look at Figure 4; its UI is much better and easier to understand compared to the Adept manager that Kubuntu likes to bundle by default.

Let me spare you the torture

As for the final words, Pardus made me use KDE3 almost after a year—and I hate the fact of downgrading to something. That said, Pardus is a pretty solid distro, which will address the requirements of most desktop users, out of the box. If you want more, the package repository probably has the answer—although, don't expect the number of packages to be as high as any of the major distros.

Pardus 2008.2

★
★
★
★
★



Pros:

- Boot speed, Tasma configuration manager, multimedia support, hardware support, impressive theme.

Cons:

- Absence of KDE4 by default, older version of OOO, spelling mistakes and missing translations

Platform: x86

Price: Free (as in beer)

Website: www.pardus.org.tr/eng



By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

Is It A Dream or What?



Dreamlinux 3.5: Debian's goodness, with a revamped user interface and out-of-the-box multimedia support.

The look and feel factor leaves a mark on everyone—as the Mac OSX has proved over the years. We have had a distro from GNU/Linux userland that is reminiscent of the OSX desktop. So, when Dreamlinux 3.5 came out this March, I decided to take it for a ride. Here I present the results of the test drive.

To begin with, here are the specs of my test system:

- AMD Phenom X3 8650 over-clocked to 3.013 GHz
- AMD 790GX motherboard with ATI HD3300 graphics
- 4GB DDR2 800 MHz RAM
- Western Digital 640GB Aaks hard disk
- HP all-in-one 1410

The onset of a beautiful dream

Dreamlinux comes in two flavours, featuring the minimal XFCE and the feature-rich GNOME. I chose the XFCE edition. This is packed into a 695 MB ISO image (which is bundled in this month's LFY DVD), while the ISO image for the GNOME edition comes packed in an awkward 913 MB size, which will push you to burn a

DVD, leaving you with a lot of wasted storage space. I really wished for a bigger ISO image if they wanted to force users to burn a DVD. Anyway, I downloaded and burned both.

Restarting the PC with the XFCE CD inside, I was confronted with a bluish boot splash screen. Now if you have used an older version of Dreamlinux, you won't be pleased, as there is no change in the boot loader screen. There are only two entries here—one to boot the XFCE desktop and the other for memory diagnostics.

I booted the disc and to my surprise it was all verbose mode. Thankfully, it didn't take too long to load the desktop—back to a GUI, thankfully!



Tips: The live CD should load the desktop without any human intervention. However, if you do run into trouble, and the system asks for a user name and password 'which is quite irrelevant for a live CD', don't panic—'dreamer' is the user name and 'dream' is the password. If you run into X server problems (for example, if the GUI fails to load), try reconfiguring your X by using the following command: `sudo dpkg-`



Figure 1: The default desktop

reconfigure xserver-xorg. Then log in by typing: startx.

The default desktop looks a bit cluttered with too many icons, but also includes a cool dock allusive of Mac OSX (Figure 1). The 'Engage' dock doesn't require Compiz to work; it uses fake transparency, just like screenlets and gDesklets. Likewise, it won't allow maximised applications to cover it up, which is sort of irritating. Shifting to Compiz will trigger the Avant Window Navigator, a much more feature-rich dock.

The wallpaper is simple and stylish, listing the distro's peers. I really liked the way it listed 'Powered by Debian' and 'Boosted by Flexiboot'. Flexiboot aims to provide a clean boot experience, but I still don't find it as snappy as some of the other new distros that boot in 20 secs or less.

On my system, the distro came up with a decent resolution of 1280 X 1024 pixels. This is impressive considering that the distro doesn't come with the proprietary ATI *fglrx* drivers.

Coming to the installation steps, the desktop lists many install icons that will surely baffle a new user. Dreamlinux is perhaps the only one of its kind that provides so many options:

1. *OEM Installer*: is aimed at an OEM installer, or someone who wants to dedicate the whole PC to Dreamlinux.
2. *Live Dream*: This install option will dump the ISO in a USB drive, making it accessible for booting as a live medium. Please note the changes you made in the live session will not be sustained after you restart the PC.
3. *Persistent Dream*: This option is the same as *Live Dream*. However, it lets you store your settings and changes, so that they're available even after a reboot.
4. *DL Installer*: Now this is the option that lets you install the distro on the hard disk. No frills here. Dreamlinux offers the easiest install wizard I have ever seen -- just one window, with no next tab to click. All options merge perfectly into one window.
5. *Portable Dream*: This option installs the whole Dreamlinux into the USB, making it a portable operating system. You can thus change things and save files in it, and you won't lose settings after rebooting the desktop. Quite a handy tool for people on the move, and with the plummeting prices of USB drives, it makes sense to have a

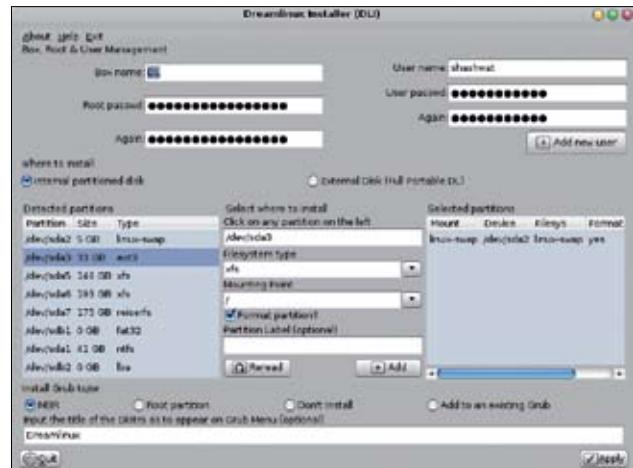


Figure 2: DL Installer



Figure 3: Compiz on the go

USB desktop in your pocket. There isn't a separate installer for this; it is integrated with DL install. Just plug in the external drive, and select the 'External Disk' option under the 'Where to Install' section of the DL Installer.

6. *MkDistro Easy Remaster*: This is a good option that allows you to create your own modified ISO with no hassle. Simply install what you want in *Live* mode and make a different ISO using the Easy Remaster.

There is another icon on the desktop named *Easy Install*, which is not for OS installation. *Easy Install* enables a one-click install for pre-selected applications.

Finally, after fiddling around a bit, I thought of installing the distro and was stunned to see that the installation finished in a mere five minutes.

Riding the dream

Post installation, I rebooted the system and noticed that the bootsplash is visible, though it goes into verbose mode now and then. Still, the transition to verbose and bootsplash is very smooth. The splash becomes faint before fading into verbose mode or switching to the log-in screen.

After logging in, I was pleasantly startled by the welcome sound, which seemed refreshingly different from the other gloomy sounds that we generally hear as the log in greeting. The desktop remains the same—the Engage Dock and the



Figure 4: Easy Install



Figure 5: The DreamLinux Control Panel

install icons were still there. I cleaned my desktop removing the install icons and installed ATI proprietary drivers... Dreamlinux Control Panel lists an install option for NVIDIA drivers, but when you click on it, you're only guided on how to install it. To install drivers, simply press **CTRL+ALT+F1** and type **nvidia-install** or **ati-install** to install the respective proprietary drivers.

I noted that the wizard was about to install older drivers. I closed the install set-up and installed **fglrx** from the latest ATI Catalyst 9.3. After rebooting, I was able to get the full resolution for my monitor. But I wasn't able to get hold of Compiz' composite effects. With some searching and help from developers, I was able to use Compiz. To enable it in the ATI/NVIDIA card, in case it doesn't work,



Figure 6: Big Buck Bunny in XBMC

open the terminal and take the following steps.

```
sudo gedit /etc/X11/xorg.conf
```

Insert the line in the 'Extensions' section, if it is not present there:

Section "Extensions"

Option "Composite" "Enable"

EndSection

Dreamlinux comes only with a handful of Compiz plug-ins, but you can install restricted and additional plug-ins—it's Debian, so you have the `apt-get` magic. You can then enable it from CCSM. Though in my case, whenever I select additionally installed plug-ins, it gets deselected automatically! Pretty weird, right?

The next thing was to set up my printer. I plugged it in and opened the DL Control Panel. Since there was a Printer Wizard, I thought of setting it up, but was unable to. Then I installed **hp-lijp**, **sane**, **sane-utils** packages, and headed to configure it again. It was detected and set up correctly. I was able to print documents, but somehow Xsane didn't recognise my printer!

Lots of home-made goodies!

On the applications front, you might be a little upset at first, but once you use all that's available you will be pleased—it still has all the necessary tools for daily work. Some of the major numbers that we all like to see are:

- Linux Kernel 2.6.28.5
- Xorg 1.3 (it's too old I know; 1.6 is out now!)
- Alsa 1.0.16 (another old release)
- Xfce 4.4.2
- GIMP 2.6.0, Inkscape 0.46
- OpenOffice.org 2.4.1 (old again!)
- Firefox 3.0.6

Xfce 4.6 was not included, because it was released only a day before Dreamlinux 3.5 came out! Refer to Table 1 for a list of applications this distro comes with, by default.

DL ships a lot of multimedia utilities—Avidemux for video editing, EasyTAG for ID3 tag editing, Totem for video and

List of Apps	
Category	Applications
Accessories	Bulk Rename, Calculator, Colour Folders, Conky, HJSplitLX, Screenshot, Text Editor, Xpad
Graphics	Cheese, Get-Photos, GIMP, Inkscape, gThumb, XSane
Network/Internet	Check-Gmail, Firefox, gFTP, Pidgin, Thunderbird
Office	OpenOffice.org suite
Multimedia	AviDemux, Brasero, EasyTAG, Totem Media Player, Rhythmbox, Sound Converter, Sound Juicer, Sound Recorder, SPDIF Switch
System tools	Bluetooth Analyser, Dreamlinux Control Panel, Easy Install, Engage Admin, Install options, Process Viewer (htop)

Table 1

Rhythmbox for audio. I didn't have a problem playing any media files. I was even able to play my Blu-ray movies in their full glory, without any hiccups. The only problem was with playing midi and dat files. For your optical disc burning need, Brasero is installed.

In addition to basic applications, DL splices up the distro by adding its own home-made Ruby scripts for pumping up the experience. There are many scripts with easy-to-use GUIs to help you set up your system. Conky lovers get a pre-configured script to set up Conky, which is pretty handy. There is a Colour Folder script that changes colours with simple clicks.

For audio buffs there's an SPDIF redirecting switch. If you have a Dolby set-up you can easily use it with this SPDIF script. The icing on the cake is the Dreamlinux Easy Install and Control Panel. Easy Install acts as a *wget* wrapper, with few selected applications on board; all you need is to click on the application to install it (Figure 4). The only let down is that it installs older versions of software. Developers should have updated the easy install repositories with newer software. But you can expect some updates in the near future.

The Control Panel serves as a centralised hub to control major parts of the system (Figure 5). It consists of basic features like setting up a network, sharing folders, enabling composite windows, etc. A few exquisite features are AMD Cool 'n' Quite, Multiscreen set-up, etc. AMD CnC works only with the AMD K8/K10 series (or above) processors. This

is an energy efficient technology that reduces the multiplying of processors, thus lowering the speed and voltage required to operate when they are not needed, and saving quite a bit of electricity. You can always turn on this feature from the motherboard BIOS if this script fails to work.

A tool called HardInfo in the Control Panel lists your hardware, like *lshw*. HardInfo comes with an easy-to-use GUI as a bonus. It covers almost all system devices with decent information, and even lists temperatures from the system diode. Apart from this, it comes with a few benchmarking tools to test your system's power. Though there isn't much data to compare with, it's still good enough.

Since it is a Debian-based distro, you need not worry about package managers or searching for additional software. You can either install it from the vast Debian repositories, or you can grab software from TuxSoftware.com, a portal very similar to Linspire's Click-N-Run utility. One thing to note is that the default server is pretty slow. Even with my 512 KBps plan I don't get sufficiently good speed while installing applications.

You might be a bit disappointed if you plan to install deb files by double clicking on them—it opens the package using the archive manager. For some reason, gdebi is not associated with deb files. To associate deb files with gdebi, simply right click on the file and select the 'Open With Other Application' option. In the 'Use a custom command' box, paste the following: *gdebi-gtk %f*. Make sure you check the 'Use as default

for this kind of files' option.

A few applications like the XBMC media centre application, Softmaker Office, OpenOffice.org 3.0, and GNOME, can be directly installed from the binary packages provided at the download page of Dreamlinux. XBMC is a very interesting player—take a look at Figure 6.

Time to get up!

Well, every distro has its ups and downs and so does Dreamlinux. The only problem I am facing is with my non-functioning scanner. Another problem that you are going to face is that the Synaptic package manager is not present in the applications menu. So, the only way to access it is by clicking the big *Apt-Get* icon on the *Engage* dock.

Anyway, the bottom line is: Dreamlinux is a distro with great potential! I really don't see a reason for not trying it. 

Dreamlinux

3.5



Pros:

Fast, uses low resources, easy to use, multimedia support, centralised Control Panel, Install Options..

Cons:

Older software, No applications in menus, shabby hardware support

Platform:

x86

Price:

Free (as in beer)

Website:

www.dreamlinux.org.br

Resources

- Home Page: <http://www.dreamlinux.com.br/>
- Release Notes: <http://www.dreamlinux.com.br/desktopedition.html>
- Forums: <http://www.dreamlinuxforums.org/>

By: Shashwat Pant

The author is a FOSS/hardware enthusiast who likes to review software and tweak his hardware for optimum performance. He's interested in Python, in Qt programming and is fond of benchmarking the latest FOSS distros.



A Distro Hopper's Final Destination?

SimplyMEPIS, just like PCLinuxOS, once had an edge over the mainstream distros, thanks to its being geared towards desktop and ease-of-use factors. Does it still hold that edge? We take a look at the latest version 8 and find out.

If you are like me, choosing the 'Right One' is always difficult. The tummy bulge of applications is despised at first sight. Beautiful looks get a second and third greedy gaze. Having the 'right asset' among applications for daily work is just too good to walk away from. A few distros like LinuxMint, DreamLinux and PCLinuxOS do have these qualities, but it is SimplyMEPIS 8 that takes the crown for this review.

It was on February 24, 2009, that SimplyMEPIS 8 was finally out after a long wait of more than a year. Built on Debian 5.0 'Lenny', you can certainly expect SimplyMEPIS 8 to be rock solid on the stability front. In fact, this is the first release from Warren Woodford, founder and lead developer of MEPIS Linux, along with the MEPISlovers community, which makes this version 8 very special. It has community inputs to the desktop theme and user manual, an auto-running intro on the CD, and some updates to the 'assistants'.

Where to get it

SimplyMEPIS8 is a live/install CD available for 32-bit and 64-bit PCs. Well, the good news is that this month's *LFY* has included the 32-bit ISO in the accompanying DVD. Select the relevant image from *distros_iso* directory of the DVD and burn it into a blank CD.

A quick look at the features

- KDE 3.5.10
- Linux 2.6.27-1-mepis-smp
- OpenOffice.org 3.0
- Firefox 3.0.6
- Guarddog firewall configuration utility
- BIND 9.6 and IPv6, enabled out-of-the-box
- Virtualisation is easily supported by downloading KVM 84 and libvirt 0.6.0 from the MEPIS 8.0 package pool
- An excellent set of MEPIS Tools similar to Mandriva Control Center (MCC) or openSUSE Yast

First impressions

Booting into the KDE desktop took about a minute and a half. The password for the root



Figure 1: The desktop

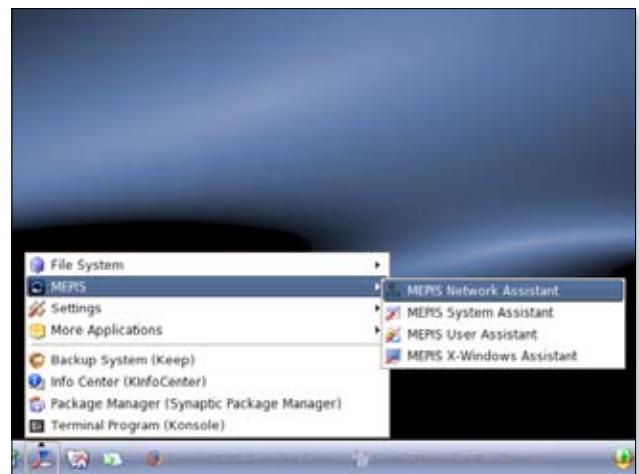


Figure 2: The MEPIS Assistants: How to trigger them

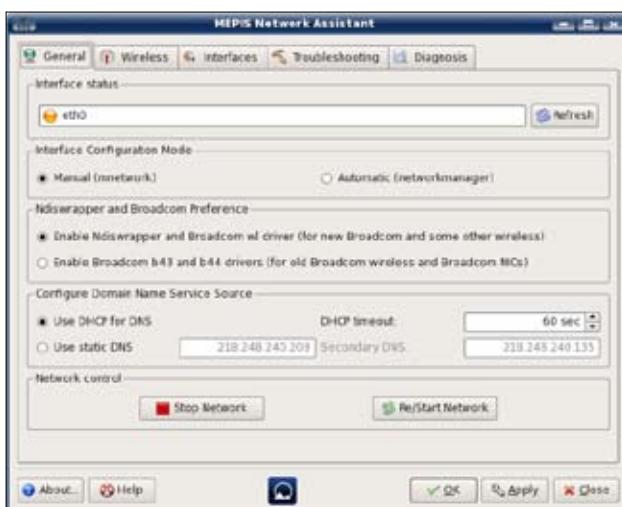


Figure 3: MEPIS Network Assistant

user is 'root' while that for the user demo is 'demo'. The cool blue water surface for the desktop wallpaper looks absolutely stunning (Figure 1). The vertically-aligned desktop icons include the MEPIS Install icon and MEPIS User Manual. The *Home* icon is present on the task bar.

Hardware detection

I have had absolutely no problem with detection of my Realtek 8139 network card, monitor resolution, HP scanner, HP Deskjet 3940 printer, USB drives, digital camera and AC 97 audio. Reading and writing to NTFS drives was a breeze. In a nutshell: everything just worked! Of course, these peripherals are fully supported under most other distros too, unless there's something royally wrong with them.

Software section

On the multimedia front, I didn't face any trouble playing DVDs with the default KMPlayer 0.10.0c. Neither did I have any problems with dropping frames (picture quality), and assorted default video settings. Of course, I would have preferred Kaffeine instead of KMPlayer. I can play MP3s without a fuss with Amarok 1.4.10—superb sound quality here too. KsCD handled all my audio CDs, and the inclusion of w32codecs was welcome. The versatile K3b 1.0.5 took care of all my burning and audio ripping needs.

Video ripping with K3b, however, required the installation of the necessary libraries, and KMPlayer refused to play VCDs and encoded

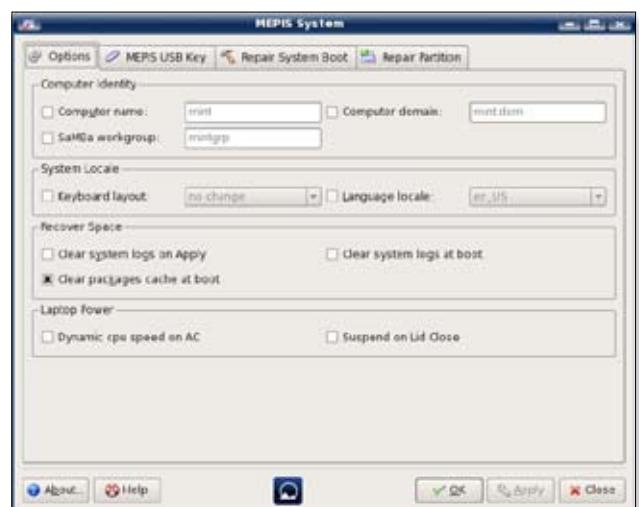


Figure 4: MEPIS Systems Assistant

DVDs. It's a no-brainer fixing these issues, but I still would have preferred it if things had worked out-of-the-box.

On the office productivity front, you have the reliable Open Office 3.0 suite, while on the graphics front, you will find the scanner utility Kooka, the digital camera utility DigiKam and a PDF reader. I wish the GIMP was also included. Yes, I sure can download missing packages like the GIMP and Kaffeine video player with Synaptic just as easily, or add the numerous packages from the Debian 5.0 DVD included with the *LFY April 09* issue to my repositories with Synaptic (*Edit→Add CDROM*).

The Internet section, by default, has the usual suspects like Firefox

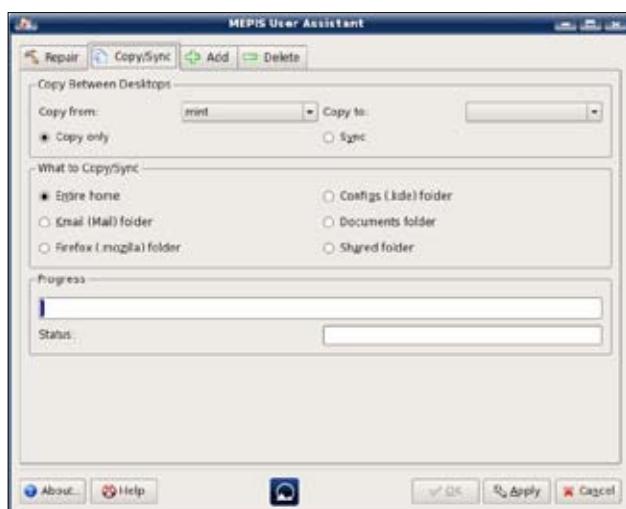


Figure 5: MEPIS User Assistant

3.0.6 and Konqueror, KMail, Guarddog (a firewall utility) and Jbidwatcher (a utility for those addicted to buying things from sites like e-bay).

MEPIS-special tools

For those who are stuck to a distro because of system-specific administration tools, viz. Mandriva Control Centre or SUSE's YaST utility, MEPIS has some special goodies in its bag. They are the MEPIS Network Assistant, MEPIS System Assistant, MEPIS User Assistant and MEPIS X-Windows Assistant. Figure 2 shows how you can access what you require.

The Network Assistant utility (Figure 3) will help you configure your wired and wireless network cards, as well as troubleshoot and diagnose them. The interface is geared towards utter simplicity. For instance, using WPA encryption for wireless networks is as simple as clicking on the WPA encryption radio button under the *Wireless* tab.

Just like Network Assistant, the MEPIS System Assistant (Figure 4) couldn't have been any easier to use. You will find tools here to repair partitions (you need to select a partition from the drop down list and click on the *Apply* button), repair system boot parameters, reinstall GRUB on the MBR/root partition and use *initrd*. There are also tools to install MEPIS on a USB key with/without encryption,

recover space from a partition by choosing to clear systems logs that accumulate over time, clear the package cache on boot, and features that enable laptop power saving options.

The MEPIS User Assistant (Figure 5) will help you add/delete users. What's interesting here is that you are able to restore Firefox config, KDE app config and group membership files in case of corruption. Even more appealing is the ability to copy files and folders between desktops! These include the entire home directory, KMail folder, Firefox (.mozilla) folder and even the .kde folder—all this by simply enabling an option. Pretty handy, eh?

Finally, the MEPIS X-Windows Assistant (Figure 6) helps you change your mouse, monitor, and graphics card (NVIDIA/ATI) settings with relative ease.

Ready to install?

I guess the MEPIS-specific tools could persuade you to install it on your hard drive, if not its Debian Lenny roots.

Installation is a painless affair that most of you will be able to manage even if it's your first time. But do take time to read the installation manual at www.mepis.org/docs/en/index.php/Mepis_8_Install if you've not played with Linux installations too often.

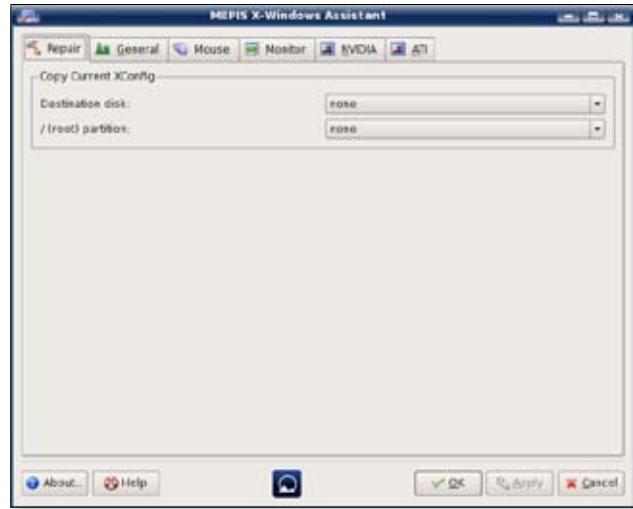


Figure 6: MEPIS X-Windows Assistant

The verdict

Whether you're a hardcore Linux user or a newbie, you just can't help falling in love the first time you use SimplyMEPIS 8. Its excellent stability, huge Debian repositories, extreme ease of use and administration, and good looks are all plus points. Indeed, the crown does go to SimplyMEPIS 8, which qualified just ahead of LinuxMint 6. Oh, wait! On the negative side, VCDs failed to play even though Kaffeine and VLC player were installed. 

SimplyMEPIS 8



Pros:

Stability of Debian Lenny, Debian admin tools, MEPIS User Assistant and Network Assistant, good multimedia support.

Cons:

VCDs don't play, takes a bit longer to login to the KDE.

Platform:

x86

Price:

Free (as in beer)

Website:

www.mepis.org

By: Nelson Lobo

The author is a lecturer at St.Xavier's College, Mapusa, Goa. He can be reached at lobonanc@gmail.com



Niyam Bhushan

Doc Till You Drop

Do you read *.doc files? Please don't.

Here's a rather pathetic joke I've just conjured for your bemused pleasure: 'Doc! Doc!' "Who's there?" "OpenOffice.org". "OpenOffice who?" "Open Your Office with ODF before you find yourself stuck with something odious."

Okay, that's pathetic even by my standards, especially in the ironic context of an article dedicated to adopting open standards, but someone needs to urgently go and bang on several doors in India. I'm talking of doors in various government offices; at large public sector undertakings (PSUs); and NGOs working in the region. As luck would have it, I've found myself professionally getting rather involved with several projects through my clients. I'm supposed to contribute creative and strategic ideas in design, publishing, and new media, i.e., until I find myself getting hit with a doc file.

One contains the 'Rfp', or the Request for Proposal. Another is the 'ToR' or the Terms of Reference. Then the usual agreement letter, an MoU, and a mandatory report. Then another, and then a few more. Heck! I'm deluged with doc files. All ordinary, simple, files filled with typed text. Nothing more. Authored in a rather expensive and proprietary word-processing software, usually bought as part of an even more expensive office-suite package. Multiply that by the number of computers in each department, of each organisation, and ditto for all the outside agencies to whom they e-mail those files.

Reloaded

Luckily enough, in the rain-drenched style of Matrix, I take on the whole blizzard of them with my extremely adept and extremely *muft* and *mukt* OpenOffice.org software. I even respond to a few with my own authored or reviewed doc files, all generated from OpenOffice.org. But, who's going to free the minds and machines of these public-funded and grants-aided organisations? Try to grasp the delicacy of the situation here. Respected and valued organisations e-mail doc files to my clients inviting them to competitively bid for projects each worth lakhs or crores of rupees. It would be quite foolish of me to exhort my clients to reject the proprietary file-formats, and insist they only receive business offers in the file-formats of their liking. "Thanks for your query on that project worth millions, but no thanks! You forgot to send me an ODF file."

"Who's going to free the minds and machines of these public-funded and grants-aided organisations? Try to grasp the delicacy of the situation here."

Living in a tea-cup

I've been living and breathing among all of you, my friends and colleagues from the free and open source (FOSS) community. We've been kicking up storms over software and open standards for several years, but alas! These are just storms in a gnu/teacup. The outside world has yet to taste the first sip of freedom. In that world, a computer is just a glorified office machine and no further thought is applied to it.

Even talking of saving a few thousand rupees in a huge, complex frame-of-things, is sub-consciously understood as a digression from the agenda at hand; and a topic to be brought up perhaps on another day, with the right people, with the right credentials, in the right context. In fact, much as my

clients are aware of the importance and value of FOSS, they find it prudent not to bring it up. I have lost count of how many meetings I've sat through, where I've been told how cash-strapped or tight the budgets are in this meltdown, while I quietly murmur to myself "Oh! Yeah? Well, then, how about getting rid of those proprietary software packages and adopting FOSS so you could pay me more, fund projects, say, three times as large, and become a financially more efficient organisation as well?"

The point here is not about offering an alternative, but making it mandatory to use FOSS and open standards, only because if it is not so in these organisations, they simply won't switch. Sunil Abraham, *mahiti.org*, has already tried hard at various international projects. I spoke extensively with Venkatesh Hariharan from Red Hat, as well as with Prakash Advani from Canonical, but until everyone reading this, pitches in with suggestions on how to stop such incredible waste of public money, and how to make digital freedom a collective goal—all we're doing is knock... knock... knocking on heaven's door. 

About the author:

Inspired by the vision of Osho. Copyright February 2009: **Niyam Bhushan**, freedomyugs at gmail dotcom. First published in LinuxForYou magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi, 'muft' means 'free-of-cost', and 'mukt' means 'with freedom.'

Run Linux on Windows

Part II

Yes, it's possible with the Ulteo Virtual Desktop.

In the first part of this article, published in the November 2008 issue of *LFY*, we looked at how to use a feature in Ubuntu that lets you install it just like a typical Windows application. Once installed, you can simply reboot your system into the world of Linux—there's no partitioning and other techie fiddling around required.

The downside there was the “reboot to use Linux”. Well, what if I tell you that there's a way to run Linux from inside Windows without a reboot? No, I'm certainly not talking about virtual machines and virtualisation here, but something as simple as a typical

click-next application installation procedure.

Allow me to introduce you to something called Ulteo Virtual Desktop [www.ulteo.com], an open source application that nicely integrates into your Windows operating system and allows you to work on a full Linux system. Its main benefit is that you can run Linux and Windows applications simultaneously within the same desktop environment without rebooting the system.

Ready, steady, go!

Before starting the installation, let's look at the hardware and software requirements. According to the website [www.ulteo.com], “Ulteo Virtual Desktop requires a x86-based PC with a modern 32-bit CPU and at least 512 MB RAM. A minimum of 4 GB of free HD space is required.” This is certainly not asking too much, I guess. My test system has the following specs:

- Pentium 4 with HT technology
- An Intel 865 motherboard
- 1 GB RAM

You can download the Ulteo VD from its website. There are two types of Ulteo products—Ulteo Application System is an installable live CD that offers a Windows alternative, and Ulteo Virtual Desktop is a coLinux-based Virtualised Ulteo workstation, which you can install on Windows. Its setup file is about 510 MB in size. It supports a full range of Linux applications, like Firefox, OpenOffice.org, KPDF, Skype, Thunderbird,

List of readily available apps

Software category Available applications

Graphics	The GIMP, Gwenview, Inkscape, Kpdf, Ksnapshot, digiKam, showFoto
Internet	Firefox, KFTPGrabber, KTorrent, Konqueror, Kopeete, Krdc, Krfb, Thunderbird, Skype, Sun Java 5.0 Web Start
Multimedia	KMix, KsCD, VLC media player, Amarok
Office	KMyMoney, OpenOffice.org (Base, Calc, Impress, Writer)
System	KBFX Configurator, KInfoCenter, KSysGuard, Keep, Konsole
Utilities	Ark, KArm, KGpg, Kjobviewer, KNotes, Kate, Skim
Settings	Appearance & Themes, Desktop, Internet & Network, KDE components, Peripherals, Security & Privacy, Sound & Multimedia, Systems Administration

Table 1



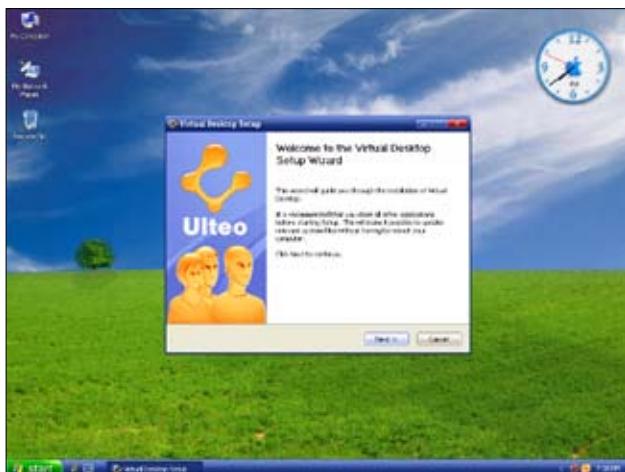


Figure 1: Virtual desktop set-up



Figure 2: The Ulteo menu for applications

The GIMP, Inkscape and many others.

After downloading it successfully, double click on the set-up (.exe) file. You will be greeted by a welcome screen (Figure 1).

Follow the on screen instructions, and if everything goes fine, within five minutes Ulteo Virtual Desktop will be installed on your Windows system. Could getting started with Linux have been any easier? It's like installing any application on Windows, with an additional benefit: it doesn't even require a Windows reboot to work properly. The virtual desktop contains a virtual filesystem of around 5 GB, as a disk file in your Windows directories where "everything will happen".

Run the Ulteo virtual desktop and you will see a panel at the top of the screen. You can browse through all the Linux applications by using the drop-down menu as shown in Figure 2. Of course, you can configure the panel according to your liking by using the *Configure Panel* option.

I guess you've figured out by now that running Linux applications alongside Windows apps at the same time is as easy as a piece of cake. Figure 3 shows some applications in action. Table 1 lists the applications included in the Ulteo virtual desktop.

The default user will be created with the user name 'me'.



Figure 3: Run Linux apps as if they are native Windows programs



Figure 4: Running a full-fledged KDE on Windows

The password for the user 'me' is, well, 'me' by default! The 'root' user password is 'root'. Ulteo Virtual Desktop contains many useful commands to work on the Linux Command Line Interface (CLI). You can also start KDE by using the *startkde* command. The KDE desktop is shown in Figure 4.

I won't carry on more about this and that feature, but request you to explore the world of Linux applications yourself. Once you get used to them, I'm sure you'll take a step in the right direction.

Of course, whenever you plan to uninstall the Ulteo VD from your Windows because you want to install an independent Linux operating system in your PC, go to the Windows control panel and under the *Add or Remove Programs* section, you will find Ulteo VD. Uninstall it the way you uninstall any other software in Windows.

However, before you do this, you better have a separate Linux installation in a dedicated hard disk partition, or else... 

By: Sandeep Kumar Yadav

Sandeep loves to test new open source tools and believes in helping people try out open source applications.

An Artist's Haven

A simple tutorial that will help your tablet digital pen to work on a GNU-powered machine.

There are many graphics designers, artists and game writers who would love to work on GNU/Linux but may find themselves helpless, because one of the most important devices that enables them to create artwork and illustrations—the digital pen—may or may not work on this operating system. Well, the problem arises because although vendors love to advertise the fact that their devices are compatible with Windows and Mac, they don't mention anything about GNU/Linux even if it's supported out of the box. So, the end-user in search of a digital pen has no clue.

I recently bought an iBall pen [WP8060]. I had checked the website and knew that it did not work on GNU/Linux. The only reason I went in for this pen was that iBall's device was the only product available in Delhi at that point of time. Despite knowing that it did not work on GNU/Linux, I deliberately asked the shopkeeper if it did. He said, "No!" But he got the message that there are GNU/Linux users interested in using this pen. I replied, "Ah, it will work—everything works on GNU/Linux."

So I got down to work and found that I might have some luck with this device after all. I must mention that the whole credit of this tutorial goes to the fine fellows who developed the drivers and other utility packages, and the bloggers who scribbled down the instructions to use these tools.

I work on Ubuntu 8.10, which will be the target platform for this tutorial. However, other distro users need not worry, as I'll have to start off by compiling from source.

First off, download the latest Wizardpen driver source [wizardpen-0.7.0-alpha2.tar.gz] from <http://tinyurl.com/wizardpen> (I have created a Tiny URL to make it easy for you. You can find the original location in the foot-note to this article.) Uncompress the package after you're done with downloading:

```
$ tar -xvf wizardpen-0.7.0-alpha2.tar.gz
```

The next step is to compile this source. However, to successfully do that, download and install the following additional packages:

- xutils

- libx11-dev
- libxext-dev
- build-essential
- xautomation
- xinput
- xserver-xorg-dev

Time to compile the driver now. Change the directory to where you extracted the wizardpen driver—for example, if you've extracted it on the ~/Desktop:

```
$ cd /home/dragon/Desktop/wizardpen-0.7.0-alpha2/
```

Now, run the following command to compile the driver:

```
$ ./configure --with-xorg-module-dir=/usr/lib/xorg/modules
$ make && sudo make install
```

After this, to check whether things have installed successfully or not:

```
$ ls /usr/lib/xorg/modules/input/wizardpen_drv.*
```

The above command should give you the following output:

```
wizardpen_drv.la
wizardpen_drv.so
```

If you see these two lines, congratulations! Your drivers are now installed! However, you're only half way through.

In order to configure your tablet, you will have to create a new *fdi* file in the */etc/hal/fdi/policy* directory. But before that we need to know the name of your Tablet pen as recognised by HAL. Plug your Tablet to your machine and then run:

```
$ grep -i name /proc/bus/input/devices
```

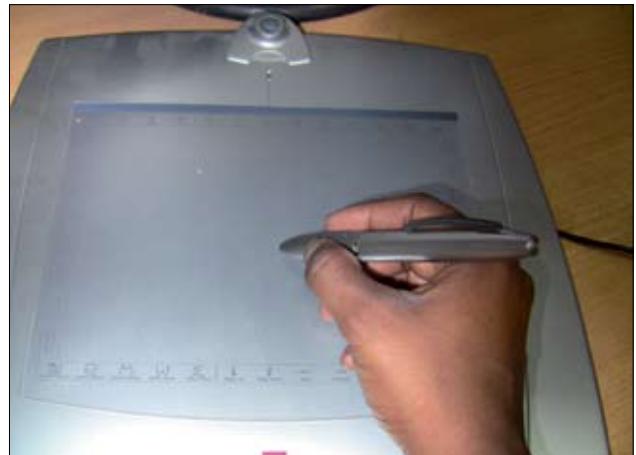
This should give you an output like this:

```
N: Name="Macintosh mouse button emulation"
N: Name="AT Translated Set 2 keyboard"
N: Name="Power Button (FF)"
N: Name="Power Button (CM)"
N: Name="PC Speaker"
N: Name="ImPS/2 Generic Wheel Mouse"
N: Name=" Tablet PF8060"
```

Note the second last line—this is the name of your Tablet as recognised by the system.

You will now have to create a policy directory for this. Create a file called */etc/hal/fdi/policy/99-x11-wizardpen.fdi* with the following content:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<deviceinfo version="0.2">
<device>
<!-- Tablet PF8060 -->
```



The iBall WP8060

```
<match key="info.product" contains="Tablet PF8060">
<merge key="input.x11_driver" type="string">wizardpen</merge>
<merge key="input.x11_options.SendCoreEvents" type="string">true</merge>
<merge key="input.x11_options.TopX" type="string">5619</merge>
<merge key="input.x11_options.TopY" type="string">6554</merge>
<merge key="input.x11_options.BottomX" type="string">29405</merge>
<merge key="input.x11_options.BottomY" type="string">29671</merge>
<merge key="input.x11_options.MaxX" type="string">29405</merge>
<merge key="input.x11_options.MaxY" type="string">29671</merge>
</match>
</device>
</deviceinfo>
```

Make sure you replace “Tablet PF8060” in line numbers 4 and 5 above with the name of your own tablet.

Time to reboot your system. When you're back, your GNU/Linux will recognise your Tablet and you should be able to use it with The GIMP to create illustrations—or just use it as a mouse.

I used to sketch a lot in the old days, and the combination of GNU/Linux and GIMP, along with this brand-new Tablet, has brought back the memories and the passion to draw. So, excuse me, as I've got to sketch some stuff.

⚠ Warning: Things do break down at times. 

Resources:

- This article is attributed to: <http://digitalbluewave.blogspot.com>
- Wizardpen driver: <http://cid-43438aff38d34c29.skydrive.live.com/self.aspx/Public/wizardpen/wizardpen-0.7.0-alpha2.tar.gz>

By: Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.



Recovering Deleted Files Easily

Here's a step-by-step guide to an application called PhotoRec, which helps you recover deleted data.

While most of us do regular back-ups of important data, some just postpone the back-up date until that dreaded day arrives—unexpectedly, in the blink of an eye, you suddenly realise all that precious data you gathered over months, has disappeared. Maybe it was just ‘the wrong key’ that got pressed. Well, it’s time to start pulling your hair out and sweating profusely. May the Lord rescue you if your boss has a bad temper.

Well don’t lose hope; there’s PhotoRec to the rescue. A huge list of 140 different file types like JPEG, MID, SQLite, Real Audio, MP3, *.doc*, Macromedia, *.exe*, *.flv*, VMware images, *.chm*, *.bz2*, Autocad, RAR, Adobe Photoshop images, etc, are supported. This tutorial will help you in recovering your data without emptying your pockets—and hey, I hope it’ll put the smile back on your face once you’re out of

your dire straits. However, you will have to spend some time searching for *that* important file.

PhotoRec is an open source multi-platform application distributed under the GPL. It is a companion program to TestDisk, an application for the recovery of lost partitions for a variety of file systems and to make disks bootable again.

Apart from Linux, PhotoRec supports the following operating systems:

- DOS/Win9x
- Windows NT 4/2000/XP/2003/Vista
- FreeBSD, NetBSD, OpenBSD
- Sun Solaris
- Mac OS X
- UNIX

PhotoRec can recover lost files from the following file systems:

- FAT/FAT32
- NTFS
- EXT2/EXT3

- HFS+
- ReiserFS (does not work very well with this file system)

Getting PhotoRec

While most distributions include TestDisk (which, in turn, has PhotoRec) in their repositories, you can download the source file or the RPM for your distro from www.cgsecurity.org/wiki/TestDisk_Download. Alternatively, you can go for PartedMagic (~ 90 MB in size), which contains TestDisk and a host of other utilities. This is available at downloads.sourceforge.net/partedmagic/pmagic-3.7.iso.zip.

The road to recovery

You can use PhotoRec to recover data or pictures that have been deleted from a pen drive. You can also recover data from a partition of a hard disk and save it to another partition on the same disk. The only condition is the partition to which data will be saved should be equal to or larger than the partition from which data will be recovered.

You will require a card reader for digital camera flash drives. Hard disks require a suitable USB enclosure. Alternatively, you can connect hard disks to an internal slot.

Step 1 : Create a directory called *photorec_dir* where PhotoRec will save files. Connect the flash/hard disk drive to your USB port/internal port (or just select your internal drive if you want to recover data from a partition) and fire up PhotoRec from the terminal as the superuser:

```
[root@localhost ~]# photorec
```

PhotoRec will display all your hard disks and USB drives. Choose the drive from which data needs to be recovered (Figure 1).

Step 2 : Up next is to choose the partition table type (Figure 2). PhotoRec supports a number of partition table types—Intel/PC, Apple, Sun Solaris, XBox, EFI GPT partition and 'None' partition types. Choose the Intel/PC type, which most of us use anyway. Even if you have a single partition, do not choose the 'None' option.

Step 3 : The next screen offers the option to recover data from the whole disk or the choice of selecting a partition. Choose your option using the up/down arrow keys. In case of a disk with multiple partitions, PhotoRec will display all the partitions, similar to what *fdisk -l* option does. Select the partition that contains the deleted data using up/down arrow keys (Figure 3).

Step 4 : We now come to the most important step of the recovery process. Select the 'File Opt' option using the left/right arrow keys and press *Enter*. PhotoRec puts forth a huge list of about 140 different

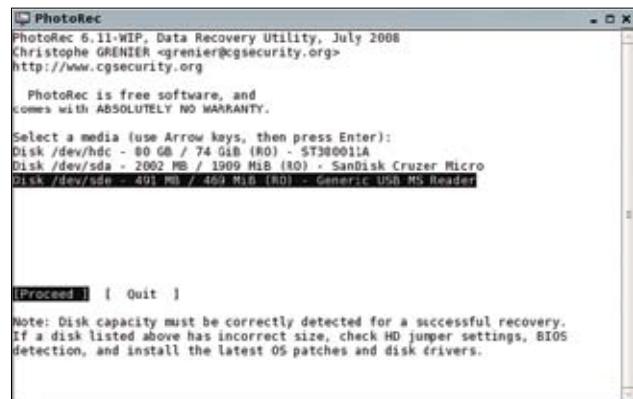


Figure 1: Choose the drive from which data has to be recovered

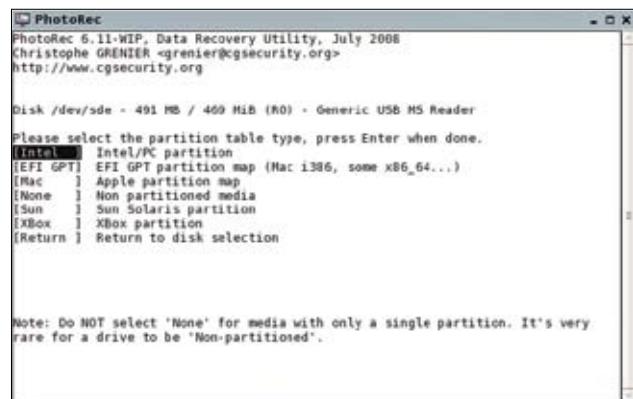


Figure 2: Choose the partition type

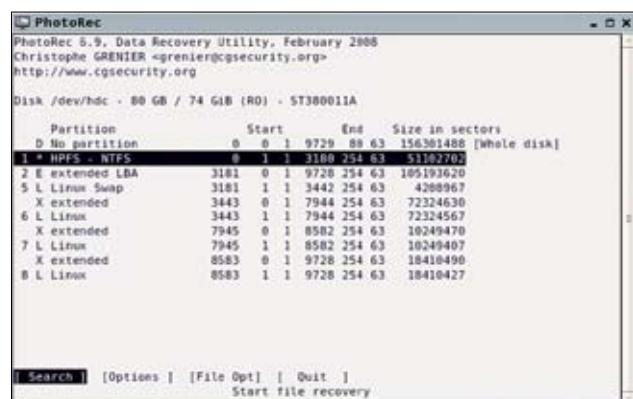


Figure 3: Choose the partition of a drive

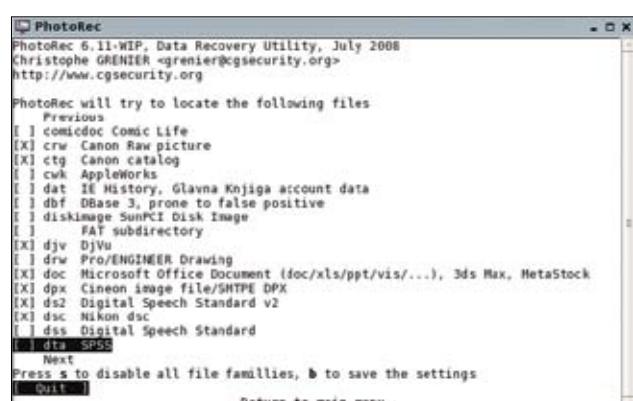


Figure 4: Choose the file types to recover



Figure 5: Choose the file system type that contains the deleted files



Figure 6: Recover files from the whole partition or from 'free space'

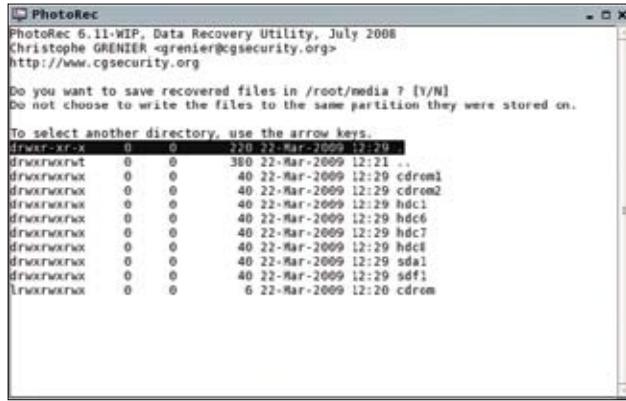


Figure 7: Choose the destination to save recovered files

file types that can be recovered. Use the up/down arrow keys to move between entries, and use the *Space* bar to select [x] or unselect [] file types to be recovered. For instance, to recover only picture files, choose 'jpg', 'gif', 'dsc', etc (Figure 4). Select 'Quit' when done. This takes you back to the last screen. Select the 'Search' option.

Step 5 : To recover lost/deleted files, PhotoRec needs to know the file system type your files were stored in. Options include 'ext2/ext3' and 'Other'. Choose 'Other' for FAT, NTFS, ReiserFS, etc (Figure 5).

Step 6 : The next screen gives the option of choosing between recovering data from the 'whole partition/disk' or from the 'free space' of the partition

(Figure 6). This 'free space' has the inodes [en.wikipedia.org/wiki/Inode] that contain the deleted data. Choose the 'free space' option using the up/down arrow keys.

Step 7 : PhotoRec now needs to know the destination folder to save the recovered files. In Step 1, we created a folder called *photorec_files*. Navigate to your path using the up/down arrow key and press *Enter*. Mine was */hdc6/home/nelson/photorec_files*. If the path is not provided, the default directory will be */root*.

Your recovered files will be saved in the destination directory in a number of folders *recup_dir.1*, *recup_dir.2*, etc. In a single *recup_dir*, you will find zip files, doc files, jpg files, etc, if you have chosen to recover these files.

Searching for a file through those that are lost can be a real pain. You need to sort these files out. Here is what you can do to sort out zip files. Make a directory for zip files as follows:

```
mkdir /home/user/Zip
```

Now as the root user:

```
mv /home/user/recup_dir.1/*.zip ~/Zip
```

Alternatively, issue the following command:

```
mv /home/user/recup_dir.*/*.zip ~/Zip
```

You can similarly repeat the steps for other file types.

Of particular interest is how to sort out picture files. Let us separate those little thumbnail pictures from your 'real' ones. Again, create a directory for small pictures with the code below:

```
mkdir ~small_jpg
```

Now, as the root user, issue the following command:

```
find /home/user/recup_dir.1/ -name "*.jpg" \
-size -20K | xargs -i mv {} /home/user/small_jpg
```

This will find all jpg files equal to or smaller than 20K and move them to *small_jpg*.

The bottom line

I found almost all of my deleted photographs and various other files with a little sweat (manually opening each file till I got the one I was really looking for) using PhotoRec. Anyone looking at a launch pad for a career in forensics? 

By: Nelson Lobo

The author is a lecturer at St.Xavier's College, Mapusa, Goa. He can be reached at lobonanc@gmail.com



New Sessions, New Tracks, New Opportunities.

ESC is where the global design engineering community gathers to learn, collaborate and celebrate innovation.

Featured Speakers:



Ben Harding
Senior Systems Engineer
SAIC



R.C. Cofer
Field Application Engineer
Avnet



Robert Oshana
Director Engineering
Freescale Semiconductors



Niall Cooling
CEO
Feabhas Limited

Who will Visit?

- Application Engineers
- ASIC Designers
- Design Engineers
- Development Engineers
- System Architects
- Hardware Engineers
- Software Engineers
- Software Associates
- Software Developers
- Software Specialists
- Software Trainees
- Engineering Management
- Corporate Management

3 Simple ways to Register:

Online: www.esc-india.com/register_now.html

SMS: <ESCI> space <Your Name> to 98337 39954

Group Booking: Call Syed at +91 22 6769 2420 or email syedj@ubmindia.com

Practical design information you can use today. Skills you use for a life time.

Learn today. Design tomorrow.



Bangalore • 22 - 24 July, 2009

**To Reserve sponsorship or booth;
speak to:**

Anees Ahmed +91 98450 32170

Shashikant Kamble +91 98336 73361

Venue: NIMHANS Convention Centre,
Near Dairy Circle, Hosur Road, Bangalore

Organised by

Produced by

Media partners

Association support





Waiting for the Android Effect!

It has been around for more than a year, but is Android set to make open source the standard for mobile phone operating systems? We take a closer look.

When it was announced with much fanfare in late 2007, many felt that Android would change the way mobile phones were used, forever. After all, it was part of the Open Handset Alliance that was dedicated to promoting open standards on mobile devices. On board were the likes of Intel, Motorola, HTC, NVIDIA and Qualcomm, who were joined a year later by ASUS and Sony Ericsson. And of course, looming above all these was the company that was the main mover behind Android—search behemoth Google. Most of the Android code has been released under the free and open source Apache licence. And the first Android phone hit the market in 2008. So where does Android stand today?

Immense potential

On the surface, there's no doubting the potential of Android. Although Linux has been used on mobile phones before, most notably by Nokia and Motorola, the devices that featured the OS did not really make headlines, barring the Motorola Ming, to an extent.

The reasons for these were many: the devices offered no discernible price advantage over others (unlike in the PC

scenario, where Linux systems have always enjoyed a significant price edge); there were never enough applications around (again, unlike the PC scenario) and the devices themselves were not the most user-friendly.

Top this off with the fact that most companies tended to 'hide' their Linux phone portfolios and highlight their proprietary or smartphone OS-driven devices, and it was no surprise that Linux had not really taken off in the mobile space, even though it had carved a niche for itself and had proved to be a stable and fast OS. In fact, not too many people know that the much-hyped Moto RAZR 2, in fact, runs on Java Linux!

In short, cell phone operating systems seemed to be dominated by the likes of Symbian, Windows Mobile, Mac OS, BlackBerry and Palm, all of which had their own issues, having imposed restrictions on application development and loaded additional cost to the price of a handset. This is why the arrival of Android was greeted by many as a major breakthrough. It was not as if people had not tried using open source or Linux-driven operating systems on phones before—it was just that never before had so many names allied themselves behind a single OS.

And Android did seem to promise a lot. There was talk of just about every sort of application, ranging from 3D games (the demo of the Android SDK showed a clip of Doom), e-mail and office suites, to just about everything. There would also be an application store (the Android Marketplace) that would allow users to download applications directly on to their phones. Best of all, Android offered the mobile user's ultimate dream—a phone that did everything at an astronomically low cost. Small wonder that many observers considered Android to be the ultimate iPhone killer.

Not all smooth sailing so far

The first Android phone, the T-Mobile G1, was released in 2008 and while it did not exactly set the mobile markets on fire, it did attract a lot of attention with its easy-to-use interface and the fact that it did come with a number of applications that other phones did not have. For one, its near-seamless integration with Google's services provided an almost desktop-like experience on the device. The fact that one could upgrade firmware over the air (without having to download software first onto a PC) as well as access dozens of applications on the Android Marketplace, led many to place it as the closest anyone had come to offering an iPhone-like touch interface. High praise indeed, considering that it cost considerably less than most devices offering similar functionality in the market.

There were, however, some rumblings. Many people felt that

the OS was too closely tied to Google for comfort and it did have certain shortcomings—one could not view mail attachments that had been sent to non-Gmail accounts, for instance! The number of applications available on the Android Marketplace, too, were limited, as compared to those found on the iPhone's App Store—you had no office suite or even a dedicated document viewer. The fact that all applications on the store were free of cost also deterred many developers from coming out with applications. The mobile space did not have a FOSS culture as yet and most professional developers expected to be paid for their effort.

Google has been working to alleviate these problems. A space was created for priced applications in February 2009, with developers getting to keep 70 per cent of the sales price and the remainder going to the carriers. The Android OS is also being considerably revamped, with the ties to Google's services being made looser, or at least less apparent, if our sources are to be believed.

The challenge: more devices, more visibility

However, the biggest challenge is that Android has been relatively absent in markets worldwide. For instance, it is altogether missing in India, one of the world's largest cell phone markets. There has also been, at the time of going to print, just one phone running the OS, which restricts user options. People might point to the iPhone's success in this regard, but then Android does not have the same marketing

muscle and advertising driving it as the iPhone does. Although a number of manufacturers—Motorola, HTC, Samsung, LG, ASUS and Sony Ericsson, to name a few—have been promising Android devices, only one has made it to the market so far. We had expected to see a number of Android devices unveiled at the Mobile World Congress at Barcelona this year but even that came to naught!

Of course, it would be very premature to write off Android. But the fact is that manufacturers need to start getting their acts together if open source is to make its presence felt in the mobile OS space. So far, even the manufacturers that are part of the Open Handset Alliance seem to be more keen on using other operating systems than the one they had all come together to support in 2007! There can be no question about Android's potential, which allied with Google's phenomenal reputation and goodwill, should make the OS a winner in the mobile space. What it needs is more devices and more visibility. And it needs them fast because, come the end of 2009 and Nokia will be unveiling the Symbian Foundation OS, another open source OS for mobile phones.

Open source is well and truly on its way to mobile phones. But will Android be its flag-bearer? The coming months will reveal all. Watch this space... 

By: Nimish Dubey

The author is a freelance writer with a passion for IT. He can be reached at nimishdubey@gmail.com

Linuxforu.com
is making a come back soon.

Automate Your Work with OpenOffice.org Macros



Here's how to automate office work with the help of macros in OpenOffice.org, with a focus on spreadsheets management.

John and Sam work in the same office. They handle a lot of documents every day and use OpenOffice.org for the purpose. Of late, John manages to finish his work early, whereas Sam has to stay in late to finish his. Moreover, John's work is more accurate and organised than Sam's. So Sam is eager to know the secret of John's success. Can you guess?

Well, John has recently learnt how to program his OpenOffice.org suite. He has prepared macros for frequent tasks, and whenever there is a need for a lot of GUI operations, he makes it simple by automating them with macros. So just with a click or shortcut key, John has managed to avoid a lot of manual and mundane repetitive tasks.

Imagine a situation where you have to create a new spreadsheet daily, and copy certain data and processes from the previous day's sheet. This will take a substantial amount of time every day. However, by preparing a macro, just a click or key combination will complete your work in a fraction of the time taken earlier.



Note: OpenOffice.org, here onwards called OOo, comes with many

components like Writer for text documents, Calc for spreadsheets, Impress for presentations, Draw for diagrams, Math for formulae and Base for database purposes. OOo documents are based on the Open Document Format (ODF) from version 2.0 onwards. The present version of OOo is 3.0.1.

Overview of macros

A macro is like a script that can avoid a series of manual operations involved in GUI applications. In OOo, macros can be coded in many languages. In this article we consider one simple language—OOBasic (OpenOffice.org Basic) to develop our macros. It is a dialect of traditional BASIC.

Some flavours of OOo come with the support of VBA macros to make documents prepared in Microsoft Office interoperable with OOo. Go-oo, a version [<http://go-oo.org>] of OOo, is one such example.

OOo comes with built-in tools for the development of macros in OOBasic, i.e., no external SDK or development tools are required. OOo provides a macro organiser to create and maintain, and an IDE to code, run and debug macros.

Your first macro

Start any Oo component (e.g., Calc), and access *Tools*→*Macros*→*Organize Macros*→*OpenOffice.org Basics*. This opens a macro organiser window. In the pane on the left, expand the current document's name, select the standard, and click on *New* to create a new module. Or, expand the standard, select an existing module and an available macro, and click on *Edit* to modify code.

This opens a basic IDE to write code (Figure 1).

Macros are classified based on the location they are stored in.

- My Macros: User written macros common for all documents
- OpenOffice.org Macros: A library of macros provided by Oo
- Document Specific: These are also user written but on a 'per document' basis and are available to the current document only.

The following is the code for a simple macro to greet people.

```
1 Sub Main
2     MsgBox "Hello World"
3 End Sub
```

Figure 2 shows the editor, where you can enter the above code. When this code is run, a dialogue appears with the message, "Hello World". This macro is available to the present document only. To make it available for all documents place it under 'My Macros'.

This example works for all types of Oo documents. In the upcoming examples, we will concentrate more on Calc documents because there is greater business value for spreadsheet automation.

Running and debugging

We can use different items in the standard toolbar of the Basic IDE to run/debug macros. A simple one is *Run Macro* (shortcut: F5). Other options are *compile*, *stop*, *step into*, *step over*, *step out*, *toggle breakpoint*, *watch*, etc. Macros can be created and maintained from this IDE also (Figure 2).

The following is the second example of a macro to access sheets, cells, etc, in a Calc document:

```
1 Sub Main
2     Dim oBook as Object
3     Dim curSheet as Object
4     Dim oCell as Object
5     Set oBook = ThisComponent
6     MsgBox "No. of sheets = " & oBook.Sheets.count
7     oBook.sheets(0).Name = "hello"
8     curSheet = book.Sheets.getByName("hello")
9     oCell = curSheet.getCellByPosition(0,0)
10    oCell.Value = 10
11 End Sub
```

In the above code:

- Lines 2-4 declare the objects *oBook*, *curSheet*, *oCell*

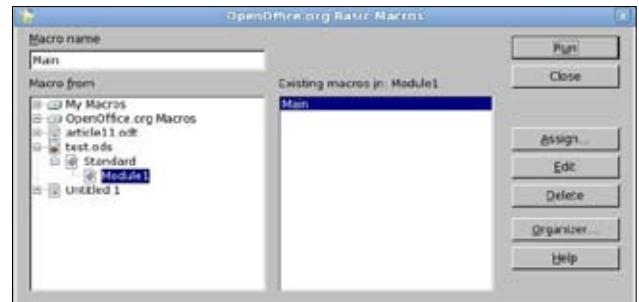


Figure 1: OOo Basic macro editor



Figure 2: The OOo macro editor

- Line 5 accesses the current document reference through *ThisComponent* and stores it to *oBook*.
- In Line 6 *oBook.Sheets* represents collection of all sheets in the present workbook and then displays the count.
- Line 7 renames first sheet to "Hello".
- Line 8 accesses the same sheet and stores it in *curSheet*.
- Line 9 accesses the A1 cell in the first sheet and stores it in *oCell*
- Line 10 changes the value of the A1 *oCell* to 10.

Let's consider some more code snippets to learn how spreadsheets can be controlled in different ways by macros.

To insert five rows before the third row in the first sheet, use the following code:

```
oSheet = ThisComponent.Sheets(0)
oSheet.rows.insertByIndex(2,5)
```

To fill range A1:A10 with a sequence, use the following line of code:

```
For i = 1 to 10
    oSheet.getCellByPosition(i-1,0)=i*
```

Next

To remove 10th column:

```
oSheet.columns.removeByIndex(9)
```

To count the number of consecutive non-empty cells in the first column:

```
n=oSheet.Columns(0).computeFunction(com.sun.star.sheet.GeneralFunction.
COUNT)
```

To search for a key pattern and highlight all cells in red:

```
Sub FindAndRed(iSheet as Object)
```

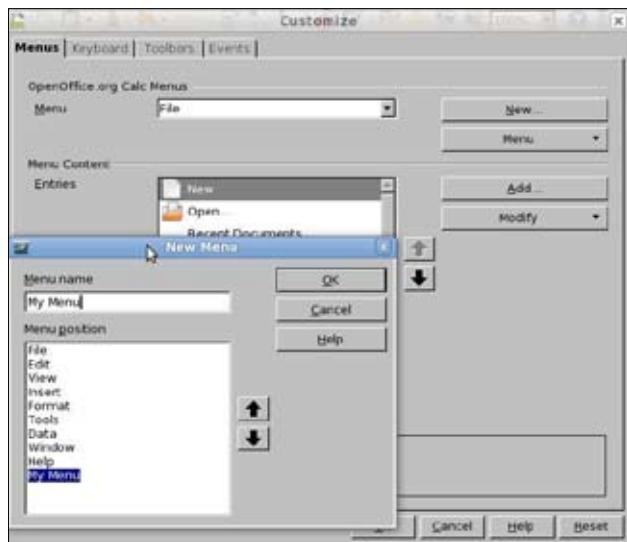


Figure 3: Creating a menu for a new macro

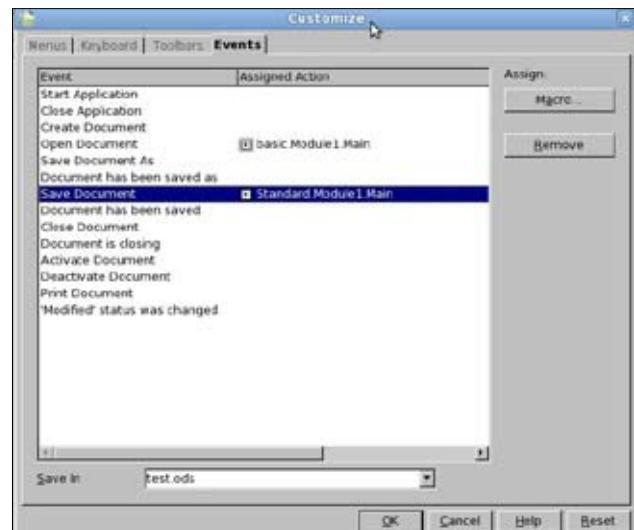


Figure 5: OOo event organiser

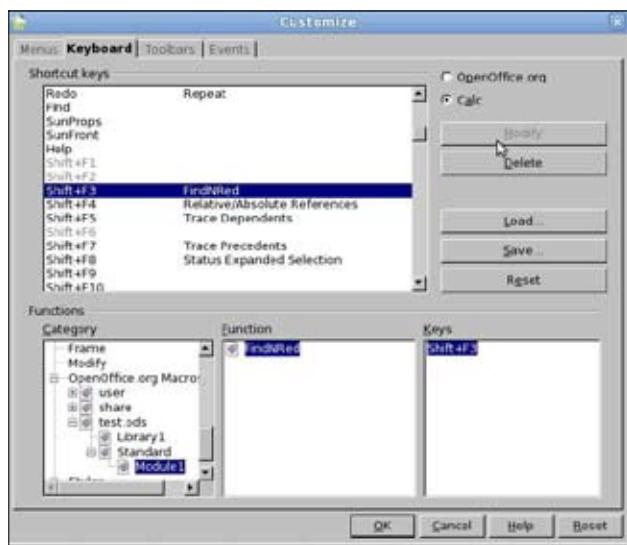


Figure 4: Creating a shortcut key

```

ThisComponent.getCurrentController.select(iSheet)
xSearchDescr = iSheet.createSearchDescriptor()
xSearchDescr.SearchString = "keypattern"
xSearchDescr.SearchCaseSensitive = true
xSearchDescr.SearchWords = true
xFound = iSheet.findFirst(xSearchDescr)
do while not IsNull(xFound)
    r=xFound.getCellAddress.Row
    c=xFound.getCellAddress.Column
    iSheet.getCellByPosition(c,r).CellBackColor =
    RGB(255,0,0)
    xFound=iSheet.findNext(xFound,xSearchDescr )
loop
End Sub

```

To repeat this process for 'N' number of sheets:

```

For i = 0 to N-1
    FindAndRed(ThisComponent.Sheets(i))

```

Next

To insert a new sheet:

Set oBook=ThisComponent

```
sNew=oBook.createInstance("com.sun.star.sheet.Spreadsheet")
```

```
oBook.Sheets.insertByName("Sheet4",sNew)
```

To create a new document:

Dim args(3) as new com.sun.star.beans.PropertyValue

```
StarDesktop.loadComponentFromURL("private:factory/scalc","_blank",0,args())
```

To apply filters for all columns in all sheets:

Dim cRange As new com.sun.star.table.CellRangeAddress

Set oBook = ThisComponent

numSheets = oBook.Sheets.Count

For i = 0 to numSheets-1

```

iSheet = oBook.Sheets(i)
numRows = iSheet.Columns(0).computeFunction(com.sun.star.
sheet.GeneralFunction.COUNT)
numCols=iSheet.Rows(0).computeFunction(com.sun.star.sheet.
GeneralFunction.COUNT)
cRange.Sheet=i
cRange.StartColumn=0
cRange.StartRow=0
cRange.EndColumn=numCols-1
cRange.EndRow=numRows-1
oBook.DatabaseRanges.addNewByName("Range"+i, cRange)
oBook.DatabaseRanges.getByName("Range"+i).AutoFilter=true

```

Next i

Easy access to macros

Macros can be accessed in a more user-friendly way:

- A new menu item can be created to invoke a macro,



Figure 6: Security manager



Figure 7: Confirm macro enabling

which can be placed under the existing menu or an altogether new menu

- A tool bar item can be created
- A keyboard shortcut can be assigned

To achieve all this, follow the options under the different tabs of the *Customize* window—*Tools*→*Customize*.

Suppose a new menu item has to be created for the previous example of *FindAndRed* (for searching a key pattern and highlighting all cells in red), then the following steps are helpful to create a menu item.

Go to *Tools*→*Customize*→*Menus*. Under *Menu List* select an existing menu or create a new one by selecting the *New* option. In this example let's create a new menu called 'My Menu' (Figure 3).

Under *Menu Content*→*Entries*, use the *Add* option. Scroll to the end and specify the macro from the available list. A new menu titled 'My Menu' will appear. Under this a menu item with the macro name of 'Main' will appear for invoking the macro. You can use the *Rename* option under the *Modify* button to change the default name from 'Main' to 'FindNRed', or anything else for that matter.

A similar process can be followed to create a toolbar item. Let's see another example of assigning a key combination for the same macro.

Go to *Tools*→*Customize*→*Keyboard*. Under the list of functions select the 'OpenOffice.org Macros' category and choose the macro. From the list of available keys select one which is free (better not to disturb standard shortcuts) and use the *Modify* option to assign it.

As you can see in Figure 4, we have assigned *Shift+F3* to the *FindAndRed* macro.

Macros as event handlers

Want to perform some task on starting or closing an application, or opening or closing a document? It's easy in OOO—put the desired code in a macro and it can be associated with some event by accessing *Tools*→*Customize*→*Events* (Figure 5).

Macro security

Although macros have their own advantages, they are dangerous too! Macros may contain vulnerable code that may harm the files and even your system. That said, there's no need to worry too much about this. OOO provides different levels of security to handle macros in documents.

- Very High: Only from trusted sources, even signed macros from untrusted sources are disabled.
- High: Signed macros from trusted sources only.
- Medium: Confirms before enabling macros while opening a document if it's from an untrusted source.
- Low: All macros will be executed without any confirmation.

To control the security level, configure *Tools*→*Options*→*OpenOffice.org*→*Security*→*Macro Security*. Refer to Figures 6 and 7.

Trusted locations and certificates can be managed from the second tab of the *Security manager* window (Figure 6).

There are a few limitations with OOBasic macros. OOBasic is a scripting language, so the code for macros is always available with the document. Even though it's a good practice as per the open source philosophy, at the user level one may damage the code accidentally (or even intentionally), which may cause a huge loss of data or corrupt the system.

Extensions written in compiled languages are a better solution in this respect, as only the binary of the extension will be available to the end user. Only trusted persons with a knowledge of programming can be provided the source code. 

References

- OpenOffice.org website: www.openoffice.org
- VBA-enabled OOO: go-oo.org
- OOO wiki—very helpful for developers: wiki.services.openoffice.org
- OOO API: api.openoffice.org

By: Rajesh Sola

The author has been involved in the OOO project since 2005 and has contributed to VBA Macro interoperability, and OOO programmability through macros and extensions. He is a faculty member of the computer science department at NBKRIST, Vidyanager. He is keen on FOSS awareness and promotion in rural areas, and is fond of teaching. He believes in training, thus encouraging and supporting students to take the open source road. You can reach him at rajesh at lisor dot org.

Fast, Efficient and Reliable Pattern Scanning and Processing

Introducing Awk, a language that has built-in features to recognise patterns and manipulate them with ease.

Whether the need is to provide some useful information to customers or the management, monitor a system's activity, extract certain lines from a file, or to automate certain tasks, a systems administrator has to often deal with text files. Carrying out these tasks usually requires parsing or manipulating text files, which might be as small as a few lines or as huge as a few gigabytes. Although there are various ways of doing this in *nix systems, most are either too trivial or so complex that one has to write several lines of code to accomplish a small task. However, Awk, with its built-in features to recognise patterns and its ability to manipulate them easily, is one of the fastest, most efficient and reliable tools. It combines the best of both worlds by providing the ease of use of grep-like tools, and the features and efficiency of a programming language. Its C-like syntax is

very easy to learn.

Awk was created by Alfred V. Aho, Peter J. Weinberger and Brian W. Kernighan (one of the creators of C language). It gets its name from the initials of its creators. Awk has evolved over the years and the current version is usually called Nawk for 'New Awk'. Gawk, the GNU version of Awk, supports the new features as well as some GNU-specific extensions. Gawk is included with the various Linux distributions and is used widely in several start-up scripts. Usually, in most Linux distributions, awk is a symbolic link to the gawk utility.

Records and fields

An Awk program divides the input file(s) into records and fields. The input file is divided into records based on a Record Separator (RS) variable. By default, records are separated by a newline character. However, an RS can be any single character or a regular expression. For example, setting RS="\$"

will separate a file into multiple records based on the occurrence of \$ in the text.

```
[testuser@localhost ~]$ cat testfile
abc$def$ghi$klm
[testuser@localhost ~]$ awk -v RS="$" '{print}' testfile
abc
def
ghi
klm
```

Output is printed based on the value of the Output Record Separator (ORS) variable. The default value of ORS is the newline character (\n). We will get into the details of this small program shortly.

Fields are, by far, the most important feature of Awk. Each input record is further divided into fields based on a Field Separator (FS) variable. Like RS, FS can be any single character or a regular expression. By default, it's a space character. If it's a space, records are separated into fields if any white space character (space, tab or newline) occurs in the text. Fields are what make Awk so useful for text manipulation. Most of the time, one has to search for a particular text at a particular location in the line. That's where fields come in handy.

Awk assigns the value of the field to a built-in variable \$n, based on the order of occurrence of that field in the record. For example, for a line containing the text "Hello, World.", if the FS value is "", \$1 becomes "Hello" and \$2 becomes "World.". These variables are valid only till the current input record. There's a special variable \$0 that is equal to the whole input record. Fields may be referenced through constants like \$1, \$2, etc, or through variables. For example if N=5, \$N may be used instead of \$5.

 **Tip:** The FS value can also be assigned using a command-line switch *-F*.

In short, each input file is divided into records based on the RS variable, and each record is further divided into fields based on the FS variable. Field values are assigned to special variables based on their occurrence, with the first field being \$1, the second \$2 and so on. The special variable \$0 is equivalent to the input record. Now, we will see how an Awk program works.

Anatomy of an Awk program

An Awk program mainly consists of patterns and actions. It can also include variable assignments and function definitions, but the most important parts are patterns and the actions that are to be taken when those patterns occur in the input text. Each pattern specified is checked against each input line read, and the actions defined for that pattern are executed. Either the pattern may be missing, in which case the defined action is executed for all input lines, or the action may

be missing—in which case the default action of printing the current input line is executed, i.e., *{print \$0}*.

The syntax of an Awk program is '*PATTERN{ACTION}*'. Action statements are enclosed within {} and the whole program is enclosed within single quotes when executed directly on the command line. An Awk program can also be saved in a file and executed using the *-f* switch. In such a case there's no need to enclose the program within quotes.

```
[testuser@localhost ~]$ cat testfile
abc,def,ghi,klm
abc,test,ghi,dsfdss
abc,def,test,kshsf
```

```
[testuser@localhost ~]$ cat myfile
/test/{ print} # pattern{ action }, Print those lines
which contain the pattern "test"
```

```
[testuser@localhost ~]$ awk -F " " -f myfile testfile
abc,test,ghi,dsfdss
abc,def,test,kshsf
```

Pattern forms

Patterns can be specified in various forms like:

- **Regular expressions:** A pattern can be any regular expression. Gawk supports extended regular expressions. Thus patterns containing character classes like [:alpha:], [:digit:], [:lower:], etc, are also supported. A detailed discussion of the regular expressions is beyond the scope of this article.
- **Relational expressions:** Relational expressions utilising the operators &&, ||, ! can be used to match complex patterns. The C ternary operator ?: is also supported for pattern matching. In this case, an expression is specified as *pattern1?pattern2:pattern3*. If *pattern1* is true, *pattern2* is evaluated, else, *pattern3* is evaluated.
- **Pattern1, Pattern2:** The man page states that this form specifies a range of text wherein the actions specified are executed for all the record lines starting with a record matching *pattern1*, continuing until a record matching *pattern2*. (See the examples at the end.)
- **BEGIN and END patterns:** There are two special patterns defined in Awk—*BEGIN* and *END*. Actions specified for *BEGIN* are executed at the start of the program before any input records are read. Thus, it's a good place for any global variable initialisation or to perform any tasks that should precede the start of the input. Similarly, actions specified for the *END* pattern are executed after all the input records have been read and the actions specified for other patterns have been executed. Actions for *BEGIN* and *END* patterns are executed only once, and are independent of the number of input records.

Actions

As mentioned earlier, action statements are enclosed within {}. They are similar to the assignment, conditional and looping statements of the C language. A statement may end with a newline character or “;”. Comments can be specified using a #.

Variables

The man page states that variables in Awk are dynamic in nature, i.e., they come into existence when they are used. Variable values are either strings or floating-point numbers. Their type is decided based on the context they are used in. They can be assigned values as *variable=value* and can be used in expressions and/or statements using their names. For example, *var=25; print var*3* will output 75. In this case, the variable *var* is treated as a floating-point number.

Besides the RS, FS and ORS variables that we covered earlier, there are a few others that are very useful while writing programs in Awk.

- **NF:** The NF variable provides the number of fields there are in the current input record. If the current record has six fields, then the value of NF would be 6. This is very useful for looping through the fields of a record when the number of fields is not known. \$NF represents the last field in a record.
- **NR:** The NR variable provides the number of input records that have so far been read.
- **FNR:** This variable provides the number of input records that have so far been read for the current input file. FNR and NR will be same if only one input file is provided. However, if there are more than one input files, FNR will show the number of records read in the current input file, whereas, NR will show the number of records read since the first record of the first input file.
- **FILENAME:** This provides the name of the current input file. If the input is specified on the command line, then its value is “-”.
- **IGNORECASE:** By default, all string operations and regular-expression matches in Awk are case-sensitive. To change that, IGNORECASE needs to be set a value greater than 0.
- **OFS:** The OFS variable specifies the field separator to be used in the output. By default, it's a space character.
- **Operators:** Operators in Awk are similar to various programming languages. Like =, +=, -= etc for assignment. Logical operators &&, ||, ! and the *increment(++)*, *decrement(--)* operators. One notable exception is the ~ operator. This operator is used to match against a regular expression. It can be negated using a ! operator to check against a negative match. Operators have a particular order of precedence. Please check the Awk man page for reference.

Control statements

Awk provides a variety of control statements that are similar to the control statements used in C. They start with keywords like *if*, *for*, *while*, *continue*, etc. This differentiates them from the simple expressions. Awk provides both the varieties of the *for* loop.

```
for (i = 1; i <= var; i++)
    print i
as well as

for (i in array)
    print array[i]
```

The second form is very useful for looping through the elements of an array, whose indices are string values. We will cover the arrays and some other advanced features of Awk in the next part of this article.

Passing variables to Gawk

Gawk supports the -v switch that allows users to pass variables to an Awk program. Shell variables can also be passed on to the Gawk program using the -v switch. The syntax is:

```
awk -v varname=${shell_var} 'pattern{ action}'
```

Once defined, *varname* can be used like any dynamic variable defined in Gawk.

Functions

There are several built-in functions available in Gawk. Users can also define their own functions and use them. We will get into the details of various functions available and how to define our own functions in the next article in this series on Awk. For now, we will take a look at a few important functions.

- **print:** As the name suggests, the *print* function allows printing of text/values. If called without any arguments, it prints the current record. Thus, *print* and *print \$0* are equivalent. Its output can be redirected to a file using *> filename*.
- **length:** The *length(str)* function returns the length of the string *str* or the length of *\$0* if no *str* is provided.
- **gsub:** The *gsub* function allows you to replace every occurrence of the text matching a pattern *ptr*, with a substitute string *str* in the target string *tr*. Its syntax is *gsub(ptr,str,tr)*. If no target string is provided, *\$0* is used. *gsub* returns the number of substitutions that took place.
- **sub:** The *sub* function is similar to *gsub*, but it replaces only the first occurrence of the pattern.

Some examples

We will make a copy of the file */etc/passwd* as *passwd2* so that we don't accidentally break the system.

- *Searching for a pattern:*

```
[testuser@localhost ~]$ pat="/sbin/nologin"
[testuser@localhost ~]$ awk -F ":" -v pat=${pat} '$NF~pat{print $1}' passwd2
```

This provides a list of all the users who are not allowed to log in to the system via the console, telnet, ssh, etc. We specify the field separator as ":" using the *-F* switch and pass the shell variable *pat* as pattern to Gawk using the *-v* switch. We match the pattern against the last field (*\$NF*) of each record and print the user-name (*\$1* is the 1st field) for the records that match the given pattern. While matching *\$NF* against the variable *pat*, we use *pat* and not *\$pat* since the latter will look for a field with the field number as *"/sbin/nologin"* which will be a null value.

- *Inserting text:* Let's suppose, we want to add a comment "Not allowed to log in" to the *passwd2* file for all users that have *"/sbin/nologin"* as shell, and "Allowed to log in" for the rest. We know that comments can be placed in the GECOS field in the *passwd* file, which is the fifth field in the file. Then, we can use the following code to add the comment:

```
awk -F ":" -v pat=${pat} -v OFS=":" '{
    if ($NF~pat) $5="Not allowed to Login";
    else $5="Allowed to login";
    print }' passwd2
```

Building on the previous example, we set *OFS*=":". The default value for *OFS* is space and the *passwd* file has ":" as the separator.

- *Replacing text:* In this example, we will try to replace */bin/bash* with */bin/sh*.

```
pat="/bin/bash"
awk -F ":" -v pat=${pat} -v OFS=":" 'BEGIN{newshell="\\"/bin/sh\""} {sub(pat,newshell,$NF); print}' passwd2
```

We utilise the sub-function to replace */bin/bash* with */bin/sh*. We also use the special *BEGIN* pattern to initialise the *newshell* variable.

- *Range pattern:* Users with UID between 500 and 510.

```
awk -F ":" -v pat=${pat} -v OFS=":" '$3~/500/, $3~/510/' passwd2
```

In this first article we covered the basic concepts of Awk. In subsequent articles we will cover advanced topics like arrays, user-defined functions, in-built functions available for string manipulation, and date functions. We will also look at some useful real-world examples like extracting a particular table from a MySQL dump, calculating the date for a few days before or after the current date, and formatting the output. 

By: Vishal Bhatia

The author is a systems manager with a VAS company, and enjoys working on Linux and open source applications. He's always keen to learn about new technologies and enjoys watching and playing cricket.



Infrastructure Box

For SMBs looking for a cost-effective but enterprise grade environment, we present the "Infrastructure Box". This comprises of Firewall, DNS, Proxy, Email, DHCP in a single physical box. This is a robust and cost-effective solution for you.

Plugin Backup

For those with an existing environment and requiring a Back-up Solution, look no further. Plug in our solution into your network and relax while it does your backup for you.

Consulting

We undertake Consulting for all IT/IS Infrastructure related requirements. Starting from the basic right to the advanced levels, we have the capability to deliver solutions that meet and exceed your expectations.

Please drop us an email at enquiry@neuronit.in to see how we can assist you further.

Neuron IT Solutions
www.neuronit.in
 Building IT Intelligently

www.eleb2b.com

Electronic Components Products Machinery Will now be sold here

Presenting Electronics Bazaar, India's First B2B Sourcing Guide For Electronics Components, Products and Machinery.

Unit Price: Rs 50
Subscription Rate: Rs 350 (12 Issues)

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase-I, New Delhi 110 020; Ph: 011-26810601-03; Fax: 26817563; Email: info@efyindia.com; Website: eleb2b.com



Building a Server from Scratch



Let's learn how to set up a Web and database server.

Yes, we are finally coming to the juicy parts. We are now ready to create our Web server, and complement it with a database (DB) server so that we can install our favourite CMS on it. So let's jump right in.

- Role: Web server
- Priority: Required
- Dependencies: Basic server, virtualisation
- Features: Two (or three) VMs, one for the Web and one for the database
- Extensible: Yes
- Ease of setting up: Very easy but needs a lot of patience
- Software set: Pound reverse proxy, MySQL/PgSQL database server, LightTPD HTTP server, interpreters for Perl, PHP and Python2
- Set-up time: Two hours

Time for some theory

Ideally we ought to take the system offline, but this is not possible since we need to fetch software from the Internet.

By now, you will have realised that with the networking set-up that we are currently using, we'll never get VMs to talk to each other. We need to change things around a bit. One solution would be to use a second network interface on all our VMs configured for internal networking. However, this is a little dirty, because there are no DNS or DHCP servers running on this interface.

The second option would be to use a Virtual Hub, or a TAP/TUN interface to bring all the VMs together, then bridge the TAP interface to eth1 (our internal interface). So all the machines in the office can access the servers. Note that we will not bridge it to eth0. *What? Then no one can access our Web server!* Relax. We will set up a reverse proxy on the server to forward all Web requests to the Web server. Our database server need not be accessible to the outside world.

Our main server will use port forwarding (between the PC and the VM, so I can give you clear instructions this time, don't worry). Any other service, such as MediaStreaming, Jabber,

etc, can use port forwarding, when necessary. Note, that we will not have a separate VM for FTP but will run FTP on our main server.

Part 1: Web server

Let's set up our Web server first. I assume that you have read through the VirtualBox manual and have mastered at least one interface. First of all, create a 500 GB (that should be enough to start with) VDI and attach it up to the Web Server VM. Start it up, create a single ReiserFS partition, and mount it under */Web*. (Do not use the mount command manually, edit *fstab* by hand and then execute *mount -a*.) Or you can use the appropriate system configuration tool.

Now, launch the *Add/Remove Programs* tool from the *Applications* menu in GNOME, and install LightTPD, php5-cgi, Perl, PHP5-MySQL and Python.

You can't, right? That's because LightTPD is nowhere in the CentOS repos. You need to add the Dag Wieers repos first. You can install the Dag Wieers repos by adding RPMForge. The instructions, though simple, involve innumerable steps and can be found at the following URL: wiki.centos.org/AdditionalResources/Repositories/RPMForge

Now install everything. After you're done, you'll need to edit the */etc/lighttpd.conf* file. Set the document root to */Web/Internet* (I hope you had *chown*'ed and *chmod*'ed it properly) and check if the handler extensions for php, pl, cgi and py are registered properly. Also check if the Python, Perl and PHP interpreters are properly referenced to use CGI (rather than FastCGI). There should be something like the following snippet in your config file:

```
cgi.assign = ( ".php" => "/usr/bin/php-cgi", ".pl" =>
"/usr/bin/perl",
".cgi" => "/usr/bin/perl", ".py" => "/usr/bin/env/ python" )
```

You'll also need to add *index.php*, *index.py*, *index.pl* and *index.cgi* to the list of default index documents, but be

wary of the order you put them in.

Please check the paths of all the interpreters. And last, but not the least, change the default port binding (*server.port*) to 8888.

Now to test your server. Create a page with the following text and save it as */Web/Internet/default.php*:

```
<? phpinfo(); ?>
```

Start up Lighty:

```
/etc/init.d/lighttpd start
```

And point your browser to *http://localhost:8888* (or *http://web.officenetwork.local:8888* if you are accessing it from outside the VM). If all went well, you should be greeted by a spick and span PHPINFO test page. If it didn't, check and re-check your config file, and then do it all over again, before trying to re-start Lighty.

There is a second part to this. All that happened earlier was to serve the public. How about our own office? We need an Intranet, with an ECMS, an Intranet mail server (somewhat like Exchange), a document collaboration system (like Google Docs), Wikis and what not. And it wouldn't be nice if someone from Tunisia, looking for his car insurance, clicked on your link from Google Search and found himself staring at the log-in page of your corporate mail server. He might try some stuff, too. So we need a separate instance of the server (a virtual server), running on our Web server, to hold our Intranet.

To start with, open up your */Web* directory and create another directory called *Intranet*. Then, copy the following snippet to your Lighty config file:

```
$SERVER["socket"] == ":80" {
    server.document-root      = "/Web/Intranet"
}
```

We are done. Just restart LightTPD, copy your *phpinfo* page over to */Web/Intranet*, and access your Intranet server.

If case you are mystified by my choice of ports, let me clear a part of your doubts. We are going to access our

p00r **JOB MARKET ?**
Campus Placement ?

Future is OPEN

Now, time to

Upgrade your Skills. Register...



RHCE

Free Training for RHCE Exam
RHCE Exam - 26th May 2009

First time in INDIA
LPI Certification Available now ...



LPIC 1 & LPIC 2
UCE (Ubuntu Certified Engineer)

Courses Starts from Apr 09
Register Now ...



Authorized Training Partner

Zend Certification

ZCE - International Certification for PHP from Zend Technologies - USA

Save

Upto

40%

Now

Spl Offer Register Zend GET
LAMP Pro Training (40 Hrs) FREE
Valid upto 24 Apr 2009



Future is open

Lynus Academy Pvt Ltd

Chennai - 044-25021630,

42171278, 9952966527

Pondicherry - 0413-2276233,

4202121, 9952966528

contactus@lynusacademy.com

www.lynusacademy.com

*Conditions Apply

Intranet site by just typing in the address of our VM. On the other hand, traffic to the public part of our Web server will be handled by Pound, a reverse proxy, which will filter our inbound traffic and automatically forward it to port 8888.

Our problem is this: We just cannot let our Intranet server be forwarded by our reverse proxy. If we do that, one fine day Google will decide to index our website, and it'll come up with all sorts of nasty trade secrets. *But our Web mail access was supposed to be via our Intranet site.* No, no! We cannot do that. What we can do, however, is to leave our Intranet site as it is and create a separate (a third) socket bind for our Mail server. If you want to know the sort of problems Google causes, go to the following site, which I found by searching for it in Google: www.partners.org/email

Anyway, create this bind:

```
$SERVER["socket"] == "9999" {
    server.document-root      = "/Web/Webmail"
}
```

Of course, create the directories and test the stuff with our PHPINFO tactic.

That finishes the Web server part. Now for our reverse proxy.

Part 2: Pound for Pound

We need to set up a reverse proxy (RP) to forward all HTTP requests from our server to our Web VM. We use a RP called Pound to do it.

Why do we need a RP? Putting all our VMs directly on the Internet would mean that everyone can access our Web, database and mail servers. That would be suicide. C'mon, once someone guesses the root password for our DB server, we're finished. We are not going to put anything directly on the Internet. What we'll do instead is create reverse proxies to filter wanted traffic, and also abstractise all our VMs and show ONE computer to the whole world. Do you really think www.google.com is just one server?

But before we set up our reverse proxy, we need to alter our networking set-up.

Part 2½: Virtual distributed Ethernet

Now push your CentOS VM into the background and come to your Ubuntu. Launch a terminal and execute the following:

```
# apt-get install vde2 bridge-utils
# echo "tun" >> /etc/modules
# modprobe tun
# tunctl -t tap0
# vde_switch -tap tap0 -daemon
# chmod -R a+rwx /var/run/vde.ctl
# brctl addbr br0
# ifconfig tap0 0.0.0.0 promisc
# brctl addif br0 tap0
# brctl addif br0 eth1
```

```
# ifconfig br0 10.111.111.254 netmask 255.255.255.0 up
```

We're done...almost! All we need to do now is tweak DNSMASQ a bit. Open up the DNSMASQ config file, and make a few changes:

1. Add one *except* line:
`except-interface eth1`
2. Add one *interface* line. This should not be required, but it's better to be safe:
`interface=br0`
3. And a CNAME record (You'll realise later why this is required)
`cname=webmail.officenetwork.local,web.officenetwork.local`

Re-start DNSMASQ. It'll now run exactly as we want it to.

Shut down the CentOS VM, and then in the VBox GUI, open Networking Settings. You had attached your machine to eth1; now separate it from eth1 and attach it to tap0. Save your settings, and re-launch the VM. Pretty soon, you'll see it boot up and get all its networking information from our server over the Virtual Network.

Back to Part 2: The Pound reverse proxy

Anyway, we need to get pounding unless we want to keep our website one great secret. To start with, query the eth0 IP address:

```
# ifconfig eth0
```

Now get running. Install Pound:

```
# apt-get install pound
```

That should do the trick. Pound is now installed. All we need to do is configure it. *Why is configuring something always a pain in the posterior?*

Anyway, open up the Pound configuration file (*/etc/pound/pound.cfg*). To start with, edit what is re-usable:

1. Change the user and group. For this series, I'm doing everything as the root. You can select anything as per your own security model.
2. I personally prefer Apache type logs. To get them, change LogLevel to 3.

Now we need to set up the actual reverse proxy part. Delete the ListenHTTP segment at the end of the file, and start from scratch. I'll give you the entire configuration in the form of a snippet as below, but you need to change the domain names as per your own domain registrations. Specifically, I'm using *sample.co.cc*, and *officenetwork.local* but you need to change the sample to your own registered domain.

Also, you got the eth0 address earlier using the ifconfig command. It should have been 192.168.1.2, but it could be anything. Just take note of that too. And make sure that the address is different for eth0 and eth1.

Use the following snippet:

```

ListenHTTP
  Address 192.168.1.2
  Port 80
  Service
    HeadRequire "Host: *www.sample.
.co.cc"
    Address web.officenetwork.local
    Port 8888
  End
  Service
    HeadRequire "Host: *webmail.sample.
.co.cc"
    Address web.officenetwork.local
    Port 9999
  End
End

ListenHTTP
  Address server.officenetwork.local
  Port 80
  xHTTP 0
  Service
    HeadRequire "Host: *webmail.
officenetwork.local"
    Address web.officenetwork.local
    Port 9999
  End
End

```

That should do it. Now you'll have to dive in to your FreeDNS account or tell your domain registrar to add a subdomain *webmail.sample.co.cc* pointing to the master site (*sample.co.cc*). Of course, use your own domain names here.

Part 3: The database

What good is a Web server without a database back-end? When you are using a database, you'll have to differentiate between your needs for the database and choose your software accordingly. By principle, MySQL is the best choice when it comes to serving as a back-end for your CMS, but is the worst choice if you are into transaction processing.

I'd recommend you create two new VMs with squeaky clean CentOS 5.3 on it (do not clone), name them *pgsql* and *mysql*, and install PostgreSQL and MySQL into them accordingly. Insert them into the network by plugging them into the TAP interface. I'm not elaborating on this due to space constraints.

Part 4: The true test

Start LightTPD on your Web VM and Pound on the server first.

Sit down at the console of your Web VM, visit www.wordpress.org and download WordPress. Extract the TGZ file into the root of your Intranet Web root directory (*/Web/Intranet*).

Now open Firefox and head over to *http://localhost*. Set up WordPress (it's pretty simple). When it comes to selecting a database server, use *mysql.officenetwork.local* or whatever you named your MySQL server. Go on accordingly. If the set-up succeeds, bingo!

Now start testing this set-up. Use a browser on the server, on any client, on other VMs, and finally from outside your office (set up something on the public site first). Some of it might not work, so troubleshoot accordingly. To start with, ask at the forums and Google Groups, post to LUG mailing lists, etc. You should get your answers soon, if not immediately. You can also reach me at *bg13_ina@users.sourceforge.net*. Questions and troubleshooting stuff are always welcome.

Next up?

Get three more subdomains, pop, smtp and imap registered for your domain name. We set up e-mail next month.



Acknowledgement

The VDE part was lifted off (though modified for our setup) from "Setting Up A Virtual Infrastructure" by Ajitabh Pandey, LFY Oct 2008, Pages 38-44.

By: Boudhayan Gupta

Boudhayan is a 14-year old student who suffers from an acute psychological disorder called Distromania. He owes his life to Larry Page and Sergei Brin. Apart from that, he enjoys both reading and writing, and when he is not playing with his Python :-), during most of his spare time he can be found listening to Fort Minor, or cooking.



TRAINING & CONSULTANCY

There is a Great Demand for NACE Certified LAMP Professionals

**LINUX
APACHE
MySQL
PHP**



**with HTML, CSS
JAVA SCRIPT & AJAX**

**DRUPAL
NCLP 10**

CHENNAI

**T.NAGAR
99405 42222**

**ANNA NAGAR
97909 25554**

**ADYAR
9790925550**

www.naceedu.com

Open Source for Career Management

Most human resource consulting firms face challenges in presenting the right candidates for right jobs. Monster.com's advert, "Caught in a wrong job", clearly outlines those challenges. Here is a case study that outlines how Open Source assists professionals in getting that coveted job.

Executive Track Associates (ETA) is a search and recruitment consulting firm with its headquarters in Delhi. It has offices in seven different locations in India—New Delhi, Mumbai, Chennai, Coimbatore, Hyderabad, Pune and Bangalore. Founded by Adarsh Matta, an engineer and alumnus of IIM-Ahmedabad, ETA works not only to provide services to its clients but also helps them to find the right fit. According to sources in the company, it provides executive solutions to clients including multinationals, entrepreneurial businesses, private equity firms, etc.

With its multiple locations and vast spread of verticals, associates at ETA faced a multitude of challenges:

- CV data was scattered in desktop mailboxes spread across the country.
- Classification, storage and retrieval of this data was a major effort.
- Just-in-time availability of the right CV was a challenge.
- Required information flow pertaining to placed candidates was a challenge.
- Moreover, advertisements in print and electronic media would result in a huge inflow of CVs, making the task of handling this vital information very cumbersome.

ETA decided that a central data repository of CVs would be the best option whereby searching the repository would be quick and yield useful information which could easily be converted to revenue.

Security was a major concern. This data bank would contain important confidential data and should not be accessible to anyone apart from authorised users. Moreover, most emails containing CVs could also be potential sources for viruses. They also did not want proprietary lock-in so that they would have to be dependant on a

single vendor for any change in the application.

They had evaluated various software options available, but none could properly address their concern on security and vendor lock-in, though a few promised extra functionality and features.

Discussing his thoughts with a colleague, Matta was suggested the open source way and referred to Varad Gupta, CTO of Keen & Able Computers Pvt Ltd, a Delhi-based IT-implementation firm that provides open source solutions to its clients. Gupta, an avid believer in the open source model comments, "Open source was the only option that I could suggest to Mr Matta. He wanted the most secure platforms at affordable costs. He wanted an IT setup which would work, rather than invest in an IT team that did most of the work. He wanted performance and reliability affordably. He wanted security from external threats and viruses and most importantly he wanted scalability without any hidden license costs. We suggested an open source model for his infrastructure and application requirements—Human Resource Management System (HRMS)."

In HRMS, the critical area of concern was the data inflow, and classification for quick access and retrieval when desired, for authorised users only. An end-to-end solution, HRMS comprises of the following:

- Mail server (to receive incoming mails, remove viruses and spam)
- VPN servers and clients (to enable site-to-site VPN within various ETA branch offices)
- Application portal for classifying, storage and search
- Reliability and redundancy of centralised data which should be available whenever desired by an authorised user.

HRMS can be implemented in various ways depending on the requirements. It includes functionalities such as payroll management,

recruitment management, performance record management and many other HR-related features which can be added as modules. Each of these functions does a specific task and automates some of the work that makes the task of the HR division easier.

The hardware and software used in the set-up

Here's a brief round-up of the specs:

- The mail server is powered by a Red Hat Enterprise Linux 5 (RHEL 5) using Postfix as the MTA (mail transport agent) on an Intel Xeon-based assembled server with 2 GB of RAM.
- The applications server is again powered by RHEL 5, which runs ETA's home-brewed ETA Recruitment Management System (ERMS). The ERMS relies on the MySQL database and runs on a Xeon-based assembled server with 4 GB of RAM.
- Distributed Replicated Block Devices are used for data storage and replication. For this, two Core-2-Duo-based server-class assembled PCs with 2 GB RAM each, and 500 GB hard disks (of course, more can be added later) are being used.

The purpose of DRBD is to provide data redundancy and high availability. Its main benefit is that if a server goes down or hard drive gets corrupted, the other server takes over and keeps on supplying data to the system.

Additionally, at ETA, the team has also created an iSCSI using a server board with a fast I/O subsystem. They have two servers, and through clustering they appear as one storage device (iSCSI block device) to the main portal server. Any data that is written, such as a new e-mail coming in, is stored on this device partition by the portal server.

Gupta explains, "The cost of implementing HRMS is minimal. The things required are the hardware infrastructure, RHEL licences, and configuration and maintenance charges. Apart from these, there are no additional licence or software costs, as Apache, PHP, MySQL and Perl are all open source."

Centralisation of data

Let's find out how the HRMS works and how data is stored in the central server. Gupta elaborates, "The mail is received by the mail server and is filtered for attachments. This module checks whether attachments exist and whether they are of the defined attachment type. If there are no attachments or if no attachments of the pre-defined type are detected, then the mail is delivered to the user. If an attachment of the desired type (could be multiple attachments and each attachment could be of different types—ODT, DOCX, XLS, PDF, etc) is found, it is stripped and stored in a central repository. A URL for viewing the stripped attachment is created and appended as a footer to the mail. The end result of this process is that new mail is delivered to the user, and the attachment is stored in a central repository."

Clicking on this URL will take the user to the central repository where the attachment is stored. The users can then categorise the attachment (as CVs, etc). All they now need is a data entry operator to do the mundane work of

filling in the CV details (like DOB, Experience, Age, etc) in the fields of a Web-based CV details form. Once classified, anyone who is authorised to view this centralised CV repository can access this CV. It also has a search option—using keywords. This CV can also be searched for any text that it contains using a content search option.

Resistance in the implementation

Any new application implementation always faces a lot of challenges, some technical and some resulting from cultural change. Gupta says, "The users had various requirements and issues, such as whether they could send multiple e-mails at a time, or if they wanted to, could they do things as they earlier did, etc." Moreover, he explains that technology is never a challenge as long as it is open source, because you can modify programs and applications to suit your needs rather than work in the manner the closed source vendor wants you to. And the concerns of new users can be taken care of by adequate training.

Key benefits

HRMS has advantages for both the end user and the companies. "As an end-user, you can e-mail your resume from anywhere (Gmail/Hotmail/Yahoo or your company mail ID) and it will be automatically uploaded. You save time and do not have to enter all your details on *that* Web page—simply e-mail the resume you have ready with you. In a second, your data is sent and uploaded onto the server, where authorised people can search and find you among the suitable candidates. For the organisation, centralised data saves a lot of money and energy."

Security features

So much of work has been done to make the data centralised, but how to secure it? When we talk about security in the context of the application, CVs can be viewed by anybody in their browser but only authorised log-ins can download the CVs onto their desktop. The second important thing regarding security is that all the branches are on site-to-site VPN (Virtual Private Network). The portal is not accessible via the Internet. Therefore, only people who are in the office or senior members of the management with VPN connections at home can log into the site.

The road map ahead

Every application has scope for improvement, and so is the case with this project. HRMS is in the first phase of its evolution, according to Matta. "We intend to take it to the second phase, where we want the whole system to be automated, such that whenever a mail is received by ETA, all the work, starting from scanning the mail to the mundane data entry work, is done by the application itself," he says. Now, this will be something to look forward to.



By: Abhijit Paul Choudhury

The author loves to hack on open source and is a gamer by heart. Oh, and he's part of the LFY bureau too.

Containing Linux Instances with OpenVZ



Understanding the OpenVZ way of virtualisation and getting started with it.

*V*irtualisation is going mainstream, with many predicting that it will expand rapidly in the next few years. Virtualisation is a term that can refer to many different techniques. Most often, it is just software that presents a virtual hardware on which other software can run. Virtualisation is also done at a hardware level, like in the IBM mainframes or in the latest CPUs that feature the VT or SVM technologies from Intel and AMD, respectively. Although a fully featured

virtual machine can run unmodified operating systems, there are other techniques in use that can provide special virtual machines, which are nevertheless very useful.

Performance and virtualisation

The x86 architecture is notorious for its virtualisation unfriendly nature. Explaining why this is the case requires a separate article on the subject. The only way to virtualise x86 hardware was to emulate it at the instruction level or to use methods like 'Binary Translation' and 'Binary Patching' at runtime. Well known software in this arena are QEMU, Vmware and the previously well-known Bochs. These programs emulate a full PC and can run unmodified operating systems.

The recent VT and SVM technologies provided by Intel and AMD, respectively, do away with the need to interpret/patch guest OS instruction streams. Since these recent CPUs provide hardware-level virtualisation, the virtualisation solution can trap into the host OS for any privileged operation that the guest is trying to execute.

Although running unmodified operating systems definitely has its advantages, there

Comparison of virtualisation software

Virtualisation software	Ability to run unmodified guests	Performance	Level
QEMU	Yes	Relatively slow	User level
Xen*	No	Native	Below OS and above hardware
KVM	Yes	Native, but devices are emulated	Hardware-supported virtualisation
UML	No	Near native	OS on OS

Table 1

*Xen can make use of VT/SVM technology to run unmodified operating systems. In this case, it looks exactly like KVM.

are times when you just need to run multiple instances of Linux, for example. Then why emulate the whole PC? VT and SVM technologies virtualise the CPU very well, but the various buses and the devices sitting on them need to be emulated. This hits the performance of the virtual machines.

As an example, let us take the cases of QEMU, Xen, KVM and UML. This comparison is kind of funny, since the guys who wrote these software, never wanted to end up in a table like Table 1. This is like comparing apples to oranges, but all we want to understand from this table is whether the VMM can run an unmodified operating system, at what level it runs, and how the performance is compared to natively running it.

Figures 1 and 2 explain the Xen and QEMU architecture.

Introducing OpenVZ

Let us suppose you want to run only Linux, but want to make full use of a physical server. You can run multiple instances of Linux for hosting purposes, education and in testing environments, for example. But do you have to emulate a full PC to run these virtual, multiple instances? Not really. A solution like User Mode Linux (UML) lets you run Linux on the Linux kernel, where each Linux is a separate, isolated instance.

To get a simplified view of a Linux system, let us take three crucial components that make up a system. They are: the kernel, the root filesystem, and the processes that are created as the system boots up and runs. The kernel is, of course, the core of the operating system; the root filesystem is what holds the programs and the various configuration files; and the processes are running instances of the programs created from binaries on the root file system. They are created as the system boots up and runs.

In UML, there is a host system and then there are guests. The host system has a kernel, and the root file system and its set of processes. Each guest has a kernel, a root file system and its own set of processes.

Under OpenVZ, things are a bit different. There is a single kernel and there are multiple root file systems. The guest's root file systems are directory trees under the host file system. A guest under OpenVZ is called a Virtual Environment (VE) or Virtual Private Server (VPS). Each VPS is identified using a name or a number, where VPS 0 is the host itself. Processes created by these VEs remain isolated from others. That is, if VPS 101 creates five processes and VPS 102 creates seven, they can't 'see' each other. This may sound a lot like *chroot* jails, but you must note the differences as well.

A *chroot* jail provides only filesystem isolation. The processes in a *chroot* jail still share processes, networks and other namespaces with the host. For example, if you run `ps -e` from a *chroot* jail, you still see a list of system-wide processes. If you run a socket program from the

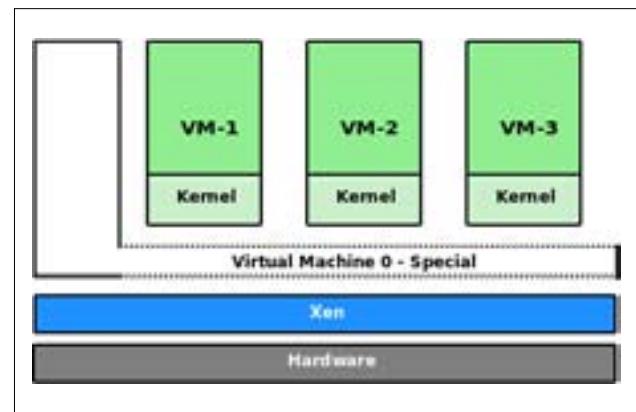


Figure 1: The Xen architecture

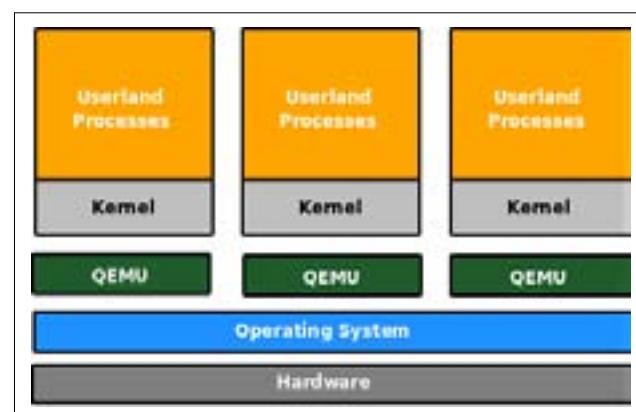


Figure 2: The QEMU architecture

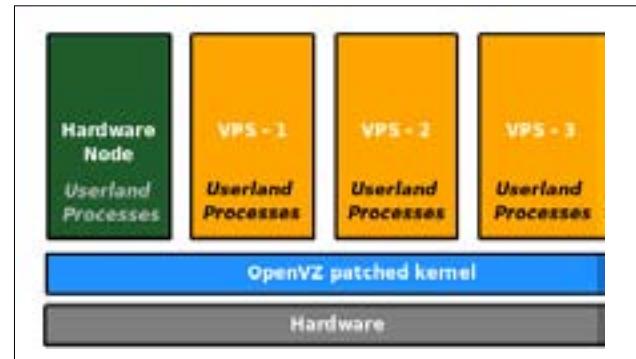


Figure 3: The OpenVZ architecture

chroot environment and listened on localhost, you can connect to it from outside the *chroot* jail. This simply means there is no isolation at the process or the network level. You can also verify this by running `netstat -a` from the *chroot* jail. You will be able to see the status of system wide networking connections.

OpenVZ is rightly called a container technology. In case of OpenVZ, there is no real virtual machine. The OpenVZ kernel is a modification of the Linux kernel that isolates namespaces and contains or separates processes created by one VPS from another (see Figure 3). By doing so, the overhead of running multiple kernels is avoided and maximum performance is obtained. In fact, the worse case overhead compared to native performance in

OpenVZ is said to be rarely more than 3 per cent. So, on a server with a few gigs of RAM, it is possible to run tens of VPSs and still have decent performance. Since there is only one kernel to deal with, memory consumption is also under check.

User bean counters

OpenVZ is not just about the isolation of processes. There are various resources on a computer system that processes compete for. These are resources like CPU, memory, disk space and at a finer level, file descriptors, sockets, locked memory pages and disk blocks, among others. At a VPS level, it is possible in OpenVZ to let the administrator set limits for each of these items so that resources can be guaranteed to VPSs and also to ensure that no VPS can misuse available resources. OpenVZ developers have chosen about 20 parameters that can be tuned for each of the VPSs.

The OpenVZ fair scheduler

Just as various resources are guaranteed to VPSs, CPU time for a VPS can also be guaranteed. It is possible to specify the minimum CPU units a VPS will receive. To make sure this happens, OpenVZ employs a two-level scheduler. The first level fair scheduler makes sure that no VPS is starved of its minimum CPU guarantee time. It basically selects which VPS has to run on the CPU next. At the scheduler level, a VPS is just a set of processes. Then, this set is passed on to the regular Linux kernel scheduler and one from the set is scheduled to run. In a VPS Web hosting environment, the hosting provider can thus guarantee the customer some minimum CPU power.

Installing OpenVZ

To install OpenVZ and have it work, you need to download or build an OpenVZ kernel, and also build or download pre-built OpenVZ tools. When you install the OpenVZ tools, it also installs the *init* scripts that take care of setting up OpenVZ. During system start-up and shut down, VEs are automatically started and shut down along with the Hardware Node (HN). Once the tools are installed, you can see that a directory named 'vz' is created in the root directory and it also contains other directories. On a production server, you may want '/vz' on a separate partition.

Installing the kernel

There are a few options here. If you are running a CentOS or RHEL distribution of Linux, pre-built kernels are available. Although the current version of Linux is 2.6.29, OpenVZ stable kernels are still in the 2.6.18 series. These are considered stable and they are either based on RHEL/CentOS kernels or the vanilla *kernel.org* 2.5.18 kernel. If you are running CentOS or RHEL, you can expect these 'stable' kernels to play well with

your system. For the sake of demonstration, I will be explaining a source-based installation.

Since we are installing from source code, first you'll need to fetch the kernel from *kernel.org*. Later on, fetch the OpenVZ patch and the OpenVZ kernel config file required to configure the kernel. Get the 'i686-SMP' config file, if you are on a 32-bit machine.

```
$ wget http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.16.18.tar.bz2
$ wget http://download.openvz.org/kernel/branches/2.6.18/028stab056.1/patches/patch-ovz028stab056.1-combined.gz
$ wget http://download.openvz.org/kernel/branches/2.6.18/028stab056.1/configs/kernel-2.6.18-i686-smp.config.ovz
```

Now, let's untar the kernel and apply the patch:

```
$ tar xvjf linux-2.6.16.18.tar.bz2
$ gunzip patch-ovz028stab056.1-combined.gz
$ cd linux-2.6.18
$ patch -p1 < ../patch-ovz028stab056.1-combined
```

Let us configure the kernel based on the OpenVZ config file and compile it:

```
$ cp ..kernel-2.6.18-i686-smp.config.ovz .config
$ make bzImage && make modules
```

Time now for installation:

```
$ sudo make modules_install
$ cp arch/i386/bzImage /boot/vmlinuz-2.6.18.openvz
$ cp .config /boot/config-2.6.18.openvz
$ cp System.map /boot/System.map-2.6.18.openvz
```

Create the initrd:

```
$ cd /boot
$ sudo mkinitramfs -o initrd.img-2.6.18.openvz 2.6.18-028stab056
```

Add the new kernel to the GRUB menu in */boot/grub/menu.lst*, and just append these lines:

```
title OpenVZ
root (hd0,0)
kernel /boot/vmlinuz-2.6.18.ovz root=/dev/sda1 ro
initrd /boot/initrd.img-2.6.18.ovz
savedefault
boot
```

Installing the OpenVZ tools

Now that the kernel installation is done, we can install the tools. The "vzctl" package contains the main utilities to manage OpenVZ virtual private servers, and the "vzquota" package contains utilities to manage disk quota. Installing these utilities is a simple affair, as

explained in the following steps. If you are on an RPM-managed system, you can download and install the RPMs (see *Resources* section at the end of this article) rather than compile from source code.

```
$ wget http://download.openvz.org/utils/vzctl/3.0.23/src/vzctl-3.0.23.tar.bz2
$ tar xjf vzctl-3.0.23.tar.bz2
$ cd vzctl-3.0.23
$ ./configure && make && sudo make install

$ wget http://download.openvz.org/utils/vzquota/3.0.12/src/vzquota-3.0.12.tar.bz2
$ tar xjf vzquota-3.0.12.tar.bz2
$ cd vzquota-3.0.12
$ make && sudo make install
```

Now that both the kernel and the utilities are installed, it's time to reboot into the new kernel. Once you reboot the system, select the "OpenVZ" option from the GRUB menu and we are booting into the system with the shiny new OpenVZ kernel. We are now ready to create Virtual Environments (VE). It is possible to create VEs based on popular distributions.

We will take the simplest approach to create VEs, which is to use a template cache. This is nothing but a file containing an archive of a Linux root filesystem. This root file system may be based on Debian, Gentoo, Ubuntu, Fedora or any distro that you prefer. There are many template caches available from the OpenVZ website. Please see the *Resources* section at the end of the article. These template caches need to be copied into the location */vz/template/cache* directory. Once they are, we can go about creating VEs based on these caches. Some template caches are available from the OpenVZ website, so let us download a few:

```
$ wget http://download.openvz.org/template/precreated/contrib/centos-5-i386-minimal-5.3-20090330.tar.gz
$ wget http://download.openvz.org/template/precreated/contrib/debian-5.0-i386-minimal.tar.gz
$ wget http://download.openvz.org/template/precreated/contrib/fedora-10-i386-default-20090318.tar.gz
$ wget http://download.openvz.org/template/precreated/contrib/ubuntu-7.10-i386-minimal.tar.gz
```

Now, let us create a VE based on Ubuntu:

```
# vzctl create 101 --os-template ubuntu-7.10-i386-minimal --config vps.basic
```

The *vzctl* command can be used for many different things in OpenVZ. In the above command line we are using it to create a VE with an ID "101". The *--os-template* option tells *vzctl* which template cache to use. It is basically the name of the template cache tar

file without the trailing *'.tar.gz'*. The *--config* option tells *vzctl* to use a base config file called 'vps.basic'. This file contains VE parameters that specify various limits and barriers among other things. There are several values in the *vps.basic* file that are applied to the VE after creating it. But these values can be overridden by specifying them explicitly for this VE. For example, we will look at how to set various values for the VE we just created, using the same *vzctl* command:

```
# vzctl set 101 --onboot --save
```

This command sets the "onboot" parameter of VE 101 to *true*. This means that the VE will be automatically started when the hardware node boots up. Each VPS has a private configuration file under */etc/vz/conf/<VE-ID>.conf*. When you specify the *--save* option to *vzctl*, the setting is applied to the VE, if it is running and is also saved to the VE's private configuration file.

```
# vzctl set 101 --hostname ubuntuve.k7computing.com --save
# vzctl set 101 --ipadd 192.168.1.101 --save
# vzctl set 101 --nameserver 192.168.1.1 --save
```

These commands set and save VPS networking parameters. It is now time to start the VPS:

```
# vzctl start 101
```

Let's view the list of VEs by running the *vzlist* command:

```
# vzlist
VPSID  NPROC STATUS IP_ADDR  HOSTNAME
101    20 running 192.168.1.101 ubuntuve.k7computing.com
```

Setting the root password and logging in via SSH

Since the VPS has an IP address, you can use SSH to log into it. But the VE must be running the SSH daemon and you must know the root password. There are multiple ways to set the root password and we shall look at two of them:

- Setting the password using the *vzctl* command

```
# vzctl set 101 --userpasswd root:test123
```

- Entering into a VE

```
# vzctl enter 101
```

```
ubuntuve# passwd
```

```
Changing password for user root.
```

```
New UNIX password:
```

```
Retype new UNIX password:
```

```
passwd: all authentication tokens updated successfully.
```

The *vzctl enter <ve-id>* command is used to "enter" into a VPS. You can then run commands as if you have logged into the VE itself. Of course, this is only possible

from the hardware node. You can even start the SSH daemon if it is not already running. Another way to start the SSH daemon or to run any other command is by using the *vzctl exec* command:

```
# vzctl exec 101 /etc/init.d/sshd start
```

This command can also be used to set the root password, by running *passwd*.

You can now SSH into this VPS from anywhere in the network. OpenVZ uses bridged networking for VEs and you can reach them over the network once they are up and running.

VPS life cycle management

Any VPS can be started, stopped or rebooted using the *vzctl* commands:

```
# vzctl start <ve-id>
# vzctl stop <ve-id>
# vzctl reboot <ve-id>
```

If you want to remove a VPS entirely from the hardware node, this can be done using the *vzctl destroy* command:

```
# vzctl destroy <ve-id>
```

This command will remove all files related to the VE. It is not possible to undo this operation.

If you run a lot of Linux servers and want to take advantage of virtualisation and get the extra performance edge, OpenVZ is the right choice for you. You can have it up and running in a few hours and there is very little learning to do. Check the *Resources* section for links to free management software for OpenVZ.



Resources

- OpenVZ home: openvz.org
- Download OpenVZ kernel RPM and patches: openvz.org/download/kernel
- OpenVZ utilities, binaries and source code: openvz.org/download/utils
- Pre-created template caches: openvz.org/download/template/cache
- Browser-based OpenVZ management tool: robovz.sourceforge.net
- GTK+ based OpenVZ management tool: easyvz.sourceforge.net

By: Shuveb Hussain

The author is head of virtualisation and cloud computing R&D at K7 Computing. He is fascinated with computers and electronics, and has learnt a lot from hackers on the Net who created the world of FOSS. He can be reached at shuveb@k7computing.com.

LINUX
For You *Calling...*
Share your
knowledge

**Inviting FOSS experts to write articles
on their area of interest**

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools



Enterprise IT 2009

The 6th International Information Technology Exhibition & Conference for the Enterprise

Asia's Unparalleled One-stop ICT Platform

EVENT HIGHLIGHTS

Exhibition

- Witness new technologies with over 175 product launches
- Meet leading suppliers from 1,300 exhibiting companies
- Network with senior level executives, trade professionals & associates

Key Exhibitors

Barracuda Networks, BlackBerry, Blue Coat Systems, DZ Card, Fortinet, Garmin, Huawei, Inmarsat, LG Electronics, NAVTEQ, NTT Docomo, PCCW Global, Samsung, SES NEW SKIES, Singapore Technologies Electronics, Yahoo!, ZTE and many others.

Hot Technologies

Learn more about Green IT, IPTV, Information Security, Mobile Entertainment, Navigation and LBS and many more

16 – 19 June 2009 | Singapore Expo

Register online before 29 May 2009 @ www.goto-enterpriseIT.com

Held in Conjunction with:



Organised by:



Singapore Exhibition Services Pte Ltd

Tel: +65 6223 6234 | Fax: +65 6223 6233
Email: events@sevents.world.com

Worldwide Associate:



Oversize Exhibition Services Ltd

Tel: +44 20 7940 2130 | Fax: +44 20 7940 2118
Email: sevents@sevents.world.com

A Part of:



Hosted by:



Endorsed by:



Held in:



Official Airline:





The Art of Guard

Part 2

Let's explore SELinux modes, and learn a few commands that will help in understanding SELinux policies better.

In Part I of this series, we discussed the basics of host security through traditional security mechanisms (DAC) and newer efficient mechanisms (MAC). We also understood that the right combination of both the Access Control Methods eventually leads to an effective security policy that can be applied to any host.

We learnt that applying security to an operating system as a whole was far more practical and efficient than applying security to each individual application running on the operating system, and that SELinux was one such mechanism. SELinux achieves MAC objectives by applying a 'Security Context' to subjects and objects, and by controlling access of subjects to objects based on access control rules.

In Part I, we also learnt to enable SELinux on Red Hat Enterprise Linux. The steps are briefly summarised below:

- Edit `/etc/sysconfig/selinux` and make sure it contains the lines:

```
SELINUX=permissive
SELINUXTYPE=targeted
```

- Reboot and check.

In this article, we will explore Permissive and Enforcing Modes. We will also learn a few SELinux commands that will help in understanding SELinux policies better.

Understanding SELinux modes

Irrespective of the policy or the rules implemented through SELinux Type Enforcement, there are three modes of operation for SELinux:

- Disabled
- Permissive
- Enforcing

Disabled mode implies that SELinux is disabled and

not implemented on the host. This has been the most common choice in installations seen by me. Hopefully, by the end of this series, we shall be able to bring about a change in that practice by encouraging more system administrators to adopt SELinux.

Permissive mode is similar to Debugging Mode. In Permissive Mode, SELinux policies and rules are applied to subjects and objects, but actions (for example, Access Control denials) are not effected. The biggest advantage of Permissive Mode is that log files and error messages are generated based on the SELinux policy implemented.

In other words, if the SELinux policy would prevent the `httpd` subject (Apache Web server) from accessing the object folder `/webdata` on my system, implementing SELinux in Permissive Mode would let the Apache Web server access the folder `/webdata` but log a denial in the log files.

This error logging informs the system administrator that if SELinux is activated in the Enforcing Mode, the `httpd` subject would be disallowed access to the `/webdata` folder on my system.

Permissive Mode is the initiating point for all those wanting to explore the world of Type Enforcement through SELinux. Without blocking access to your favourite programs such as OpenOffice.org, Evolution, etc, it provides you with enough debugging information to fine tune your policy before deploying it on your system.

Enforcing Mode, as the name signifies, is SELinux in action. All production systems, when hardened, should enable SELinux in Enforcing Mode. SELinux through Access Controls does have a minor performance overhead, but compared to the advantages that it brings to the table, I am sure it will soon become the norm to implement SELinux on production servers.

Controlling SELinux

The *getenforce* command gets the current mode of SELinux. It reports whether SELinux is Enforcing, Permissive, or Disabled. On a system with SELinux disabled, it will display the following:

```
[root@station20 ~]# getenforce
Disabled
```

On a system with SELinux in Permissive Mode, it will display:

```
[root@vbg ~]# getenforce
Permissive
```

On a system with SELinux in Enforcing Mode, the following will be displayed:

```
[root@vbg ~]# getenforce
Enforcing
```

setenforce modifies the mode SELinux is running in. It is used to toggle between Permissive and Enforcing mode when SELinux is enabled.

To activate "Enforcing mode" on an SELinux-enabled system, run:

```
[root@vbg ~]# setenforce 1
```

To check the current status, use *getenforce*. To activate Permissive SELinux mode, execute:

```
[root@vbg ~]# setenforce 0
```

The *sestatus* command is used to get the status of a system running SELinux. Apart from mentioning the current mode of SELinux, it gives more information about the SELinux policy.

On a system with SELinux disabled, it will display:

```
[root@station20 ~]# sestatus
SELinux status: disabled
```

On a System with SELinux in Permissive Mode, it will display:

```
[root@vbg ~]# sestatus
SELinux status: enabled
SELinuxfs mount: /selinux
Current mode: permissive
Mode from config file: permissive
Policy version: 21
Policy from config file: targeted
```

The first line informs us that SELinux is enabled in this system.

The second line is of great significance. It displays the mount point of the SELinux pseudo file system. This file system is quite like the proc and sys file systems, and contains run-time information about your SELinux mode and various other things.

**FOSTERing Linux is The Winner of the
'Most Promising Certified Training Partner'**
&
The 1st Institute in India to produce RHCS5s



Varad Gupta
India's 1st RHCS5



Rajveer Singh
India's 2nd RHCS5



Ashutosh Taiwal
RHCS5



Vishal Bhatia
RHCS5



Kapil Duggal
EX 333 & EX 423

Bring this coupon
&
Get 10% discount
on
any of our trainings

**The Next RHCS5
This could be YOU!!!**

Get Certified on:

- Red Hat Certified Engineer
- RHS333 - Red Hat Enterprise Security: Network Services
- RH423 - Red Hat Enterprise Directory Services & Authentication
- RHS429 - Red Hat Enterprise SELinux Policy Administration
- RH142 Linux Troubleshooting Techniques and Tools
- RH184 Red Hat Enterprise Linux Virtualization
- Clustering with Virtualization
- Mail Server
- Bash Shell Scripting



Free & Open Source Training Education & Research

SCO M-22, Second Floor Old DLF Colony,
Sector-14 Gurgaon 122 001,
Haryana, Tel: +91-124-4268187, 4080880

55, Deepak Building 13, Nehru Place
New Delhi 110 019 Tel: +91-30880047 / 49
Email Id: info@fosteringlinux.com



TRAINING PARTNER

You can change run time parameters of the SELinux system by directly writing to the files in this pseudo file system. As an example, just issue this command as the root user:

```
[root@vbg ~]# echo 1 > /selinux/enforce
```

You will see that the mode of SELinux has changed from Permissive to Enforcing.

To return back to Permissive Mode, you can run either of the commands:

```
[root@vbg ~]# echo 0 > /selinux/enforce
```

or

```
[root@vbg ~]# setenforce 0
```

The third line mentions the current SELinux mode, whereas the fourth line mentions the SELinux mode under which the system booted.

The fifth line mentions the version number of the policy (we will come to this later in this series) and finally, the sixth line mentions the Policy loaded from the configuration file (*/etc/sysconfig/selinux*) at boot time.

Understanding various types of policies

By definition, an SELinux policy is a collection of rules for SELinux Mandatory Access Controls. Each one of us can make a policy to suit our needs much like we define firewall rules through *iptables*. There can be no standard policy that can apply to all situations.

By default, there are two policies shipped along with Red Hat Enterprise Linux: Targeted and Strict.

The Targeted Policy is the first step in assisting system administrators to understand and implement SELinux. It only ‘targets’ certain network daemons such as the Apache Web server, FTP server, BIND DNS server and a few others, while leaving the vast majority of end-user applications largely untouched. It creates an ‘unconfined’ domain ‘confinement’ (interesting paradox, isn’t it?) and does not apply Access Control Restrictions to most applications in the unconfined domain.

This allows sysadmins to concentrate on the really vulnerable network applications and services while not interfering with their daily tasks.

Once the nuts and bolts of SELinux are clear to administrators, they should move forward towards implementing the SELinux Strict Policy.

The Strict Policy, on the other hand, is a true restrictive Access Control Policy. Before implementing this policy, make sure you understand SELinux concepts and policies well.

Under the hood of Targeted Policy

To understand any SELinux policy, use the informative command—*seinfo*.

```
[root@vbg ~]# seinfo
```

Statistics for policy file: /etc/selinux/targeted/policy/policy.21

Policy Version & Type: v21 (binary, MLS)

```
Classes: 61 Permissions: 220
Types: 1513 Attributes: 148
Users: 3 Roles: 6
Booleans: 210 Cond. Expr: 186
Sensitivities: 1 Categories: 1024
Allow: 82518 Neverallow: 0
Auditallow: 28 Dontaudit: 5086
Role allow: 5 Role trans: 0
Type_trans: 1398 Type_change: 17
Type_member: 0 Range_trans: 23
Constraints: 47 Validatetrans: 0
Fs_use: 15 Genfscon: 64
Portcon: 264 Netifcon: 0
Nodecon: 8 Initial SIDs: 27
```

seinfo is a policy query tool that queries policy files and provides vital information about it. When executed without any arguments, it queries the default loaded policy file.

As you can see the policy file read by the *seinfo* tool by default is */etc/selinux/targeted/policy/policy.21*. Installed by the *selinux-policy-targeted* RPM, this file contains the binary Targeted Policy.

The second line of the above output mentions that this is a binary policy with MLS (Multi Level Security). We will come to MLS in the later part of this series.

Now comes the interesting part of the output. You can clearly see various components of the policy. We will discuss these components in the course of these articles. Let us concentrate on a few important ones right now.

A typical security context as we discussed in Part 1 of this series, is of the type:

User Identity:Role:Type/Domain

As we can see from the above, in the default SELinux Targeted policy, there are:

- Users (3 in number)
- Roles (6 in number)
- Types (1,513 in number)

This means that any object or subject in the SELinux Policy installed in the system can have one of three user identities, one of six roles and one of the available 1,513 types.

To list user identities defined in the SELinux Targeted Policy, run the following command:

```
[root@vbg ~]# seinfo -u
```

```
Users: 3
system_u
root
user_u
```

Do check all possible security contexts for objects (files,

dirs, etc) and subjects (processes) in your system. The user identity component of the Security context will have one of the above three identities and nothing else.

Similarly, to check the available roles, use the following code:

```
[root@vbg ~]# seinfo -r
```

```
Roles: 6
staff_r
user_r
object_r
secadm_r
sysadm_r
system_r
```

And please check the available list of types yourself by using the command given below:

```
[root@vbg ~]# seinfo -t
```

Thus we now know that any existing object or subject in the system will have a security context created out of these three users, six roles and 1,513 types.

You can alter the security context of any object by using the *chcon* command.

To change the type of an object, use *chcon -t*. To change the user identity, use *chcon -u* and *chcon -r* for role.

To test the above, create an empty file “context” in the */tmp* directory and check its default security context:

```
[root@vbg ~]# touch /tmp/context
[root@vbg ~]# ls -lZ /tmp/context
-rw-r--r-- root root user_u:object_r:tmp_t:0 /tmp/context
```

To change the type of this object from *tmp_t* to *unconfined_t*, use the code below:

```
[root@vbg ~]# chcon -t unconfined_t /tmp/context
[root@vbg ~]# ls -lZ /tmp/context
-rw-r--r-- root root root:object_r:unconfined_t:0 /tmp/context
```

I leave it as an exercise for you to change the role and the user for this object. If you face any issues, feel free to contact me at the e-mail ID provided below.

Assuming that, as above, you have changed the security context of an object quite a few times and you would like to revert back to the original/default security context, *restorecon* comes to the rescue.

restorecon restores files to their default security contexts. The verbose option of the *restorecon* command also displays the changes made in the security context. To check the usage of this very handy command, look at the example below:

```
[root@vbg ~]# touch /home/vbg/test-context
[root@vbg ~]# ls -Z /home/vbg/test-context
```

```
-rw-rw-r-- vbg vbg
user_u:object_r:user_home_t:0 /home/vbg/test-context
```

```
[root@vbg ~]# chcon -t tmp_t /home/vbg/test-context
```

```
[root@vbg ~]# ls -Z /home/vbg/test-context
-rw-rw-r-- vbg vbg user_u:object_r:tmp_t:0 /home/vbg/test-context
```

```
[root@vbg ~]# restorecon -v /home/vbg/test-context
restorecon reset /home/vbg/test-context context user_u:object_r:tmp_t:0->user_u:object_r:user_home_t:0
```

```
[root@vbg ~]# ls -Z /home/vbg/test-context
-rw-rw-r-- vbg vbg user_u:object_r:user_home_t:0 /home/vbg/test-context
```

The above snippet shows the utility of the *restorecon* command. It reads the default contexts to be applied from the policy files and applies the default security context to a file/directory object.

A word of caution when using *restorecon*: DO NOT use the *restorecon* command with the *-r* (recursive) option. It may overwrite the security contexts of some important files in the system that you may have changed.

If you think that you have spoilt the security contexts of the files in your system beyond recovery, do not panic. Help is available in the form of Auto-Relabel, at boot. Simply create an empty file */autorelabel*. Please note that it is a hidden file.

```
[root@vbg ~]# touch /.autorelabel
[root@vbg ~]# reboot
```

Following the procedure mentioned above will cause SELinux to relabel the files on your system upon rebooting. Please use this to fix any improper security context on files and directories.

You can also use the *fixfiles* command to achieve the above. *fixfiles* can prevent you from rebooting your system but may not be as effective. Depending on the options and time available, you can choose any option that suits you—though I would suggest the reboot option.

Still to come

So, that's about it for now. I'll be back the following months with topics such as:

- Understanding the Targeted Policy - Part II
- Policy Modules
- MLS and MCS



By: Varad Gupta

Varad is an open source enthusiast who strongly believes in the open source collaborative model not only for technology but also for business. India's first RHCSA (Red Hat Certified Security Specialist), he has been involved in spreading open source through Keen & Able Computers Pvt Ltd, an open source systems integration company, and FOSTERing Linux, a FOSS training, education and research training centre. The author can be contacted at varad.gupta@fosteringlinux.com

Balancing Traffic Across Data Centres Using LVS



In the previous article in this series ('Data Centre Redundancy', page 56, April 2009) we got an overview of the architecture associated with data centre redundancy. An important aspect of this is to make sure that the traffic is evenly distributed across all the data centres in order to fully utilise the available capacity. In this concluding article, we will discuss the use of LVS (Linux Virtual Server) to load balance the traffic across data centres.

Lhe LVS (Linux Virtual Server) project was launched in 1998 and is meant to eliminate Single Point of Failures (SPOF). According to the linuxvirtualserver.org website: "LVS is a highly scalable and available server built on a cluster of real servers, with the load balancer running on Linux. The architecture of the server cluster is fully transparent to the end user, and the users interact as if it were a single high-performance virtual server. The real servers and the load balancers may be interconnected by either a high speed LAN or by a geographically dispersed WAN."

The load balancer is the single entry point into the cluster. The client connects to a single known IP address, and then inside the virtual server the load balancer redirects the incoming connections to the server(s) that actually does the work according to the scheduling algorithm chosen. The nodes of the cluster (real servers) can be transparently added/removed, providing a high level of scalability. The LVS detects node

failures on-the-fly and reconfigures the system accordingly, automatically, thus providing high availability. Theoretically, the load balancer can either run IPVS or KTCPVS techniques for load balancing, but owing to a very high stability of IPVS, it is used in almost all the implementations I have seen. See the sidebar titled "IPVS v/s KTCPVS" for a brief note on the differences between the two. IPVS provides Layer 4 load balancing and KTCPVS provides Layer 7 load balancing (see the sidebar).

There are three load balancing techniques used in IPVS:

- LVS/NAT – Virtual Server via NAT
- LVS/TUN – Virtual Server via Tunnelling
- LVS/DR – Virtual Server via Direct Routing

A brief overview of these techniques can be found in the sidebar titled 'IPVS Load Balancing Techniques'.

Since our real servers are located in two different data centres, we will be focusing on LVS/TUN.

Installing and configuring IPVS

Please note that the set-up explained here should only be used as a guideline and for an understanding of how IPVS works. Networking scenarios are different for every case and may demand extra reading and experimentation before getting a working set-up. My advice is that before trying this out in the field, make sure enough experiments have been done in the laboratory. Also, it is advisable to read through the documents in the *References* section at the end of the article.

IPVS v/s KTCPVS

IPVS or IP Virtual Server is an implementation of Layer 4 load balancing inside the Linux kernel. Layer 4 load balancing works on OSI Layer 4 (Transport Layer) and distributes requests to the servers at the transport layer without looking at the content of the packets.

KTCPVS or Kernel TCP Virtual Server is an implementation of Layer 7 load balancing in the Linux kernel. Layer 7 load balancing is also known as application-level load balancing. The load balancer parses requests in the application layer and distributes requests to servers based on the content. The scalability of Layer 7 load balancing is not high because of the overhead of parsing the content.

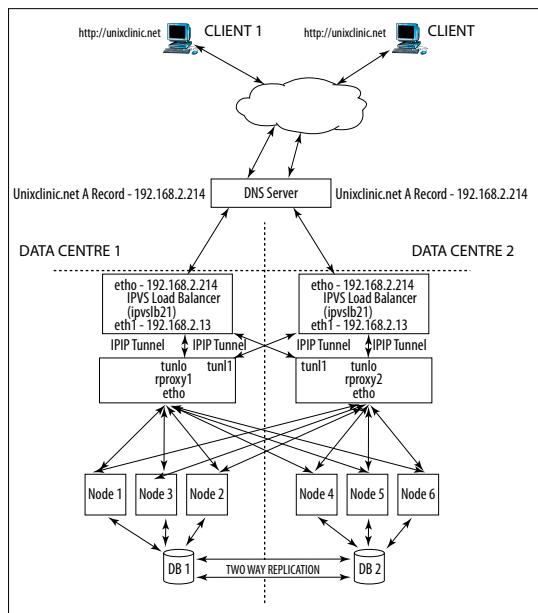


Figure 1: A diagram representation of the data centre set-up

On Debian and the likes, issue the following code:

```
# apt-get install ipvsadm keepalived
```

On Red Hat and the likes, use the following:

```
# yum install ipvsadm keepalived
```

The kernel module *ip_vs* and *ipip* may need to be loaded, but in my experience, these modules were automatically loaded when I used the *ipvsadm* command.

To start with, we will consider a scenario that has two data centres. There is one LVS load balancer in each data centre. For the sake of giving them names, we will call them *ipvs1b11* and *ipvs1b21*. Now we will configure the IPIP tunnel between the load balancers and the real servers—*rproxy1* and *rproxy2*, where *rproxy1* is in the first data centre and *rproxy2* is in the second.

Before we start the command configuration, have a look at Table 1 and Figure 1.

Please look at the following configuration carefully, and look at the host names in the prompt to identify to which server you need to type these commands:

```
rproxy1# modprobe ipip
rproxy1# ip addr add 192.168.1.214/32 dev tunl0
rproxy1# ip addr add 192.168.2.214/32 dev tunl1
rproxy1# ip link set tunl0 up arp off
rproxy1# ip link set tunl1 up arp off
```

```
rproxy2# modprobe ipip
rproxy2# ip addr add 192.168.1.214/32 dev tunl0
rproxy2# ip addr add 192.168.2.214/32 dev tunl1
rproxy2# ip link set tunl0 up arp off
rproxy2# ip link set tunl1 up arp off
```

Data Centre details				
Data Centre	Host	Interface	IP Address	Role
Data Centre 1	ipvs1b11	eth0	192.168.1.214/24	VIP
Data Centre 1	ipvs1b11	eth1	192.168.1.13/24	DIP
Data Centre 2	ipvs1b21	eth0	192.168.2.214/24	RIP
Data Centre 2	ipvs1b21	eth1	192.168.2.13/24	DIP
Data Centre 1	rproxy1	eth0	192.168.1.14/24	RIP
Data Centre 1	rproxy1	tunl0 (no ARP)	192.168.1.214/32	VIP
Data Centre 1	rproxy1	tunl1 (no ARP)	192.168.2.214/32	VIP
Data Centre 2	rproxy2	eth0	192.168.2.2/24	RIP
Data Centre 2	rproxy2	tunl0 (no ARP)	192.168.1.214/32	VIP
Data Centre 2	rproxy2	tunl1 (no ARP)	192.168.2.214/32	VIP

Table 1

Now let us add two virtual servers—one to serve normal Web traffic and one for secure traffic on Port 443:

```
ipvs1b11# ipvsadm -A -t 192.168.1.214:80 -s wlc
ipvs1b11# ipvsadm -A -t 192.168.1.214:443 -s wlc
ipvs1b11# ipvsadm -a -t 192.168.1.214:80 -r 192.168.1.14 -i
ipvs1b11# ipvsadm -a -t 192.168.1.214:443 -r 192.168.1.14 -i
ipvs1b11# ipvsadm -a -t 192.168.1.214:80 -r 192.168.2.2 -i
ipvs1b11# ipvsadm -a -t 192.168.1.214:443 -r 192.168.2.2 -i
```

```
ipvs1b21# ipvsadm -A -t 192.168.2.214:80 -s wlc
ipvs1b21# ipvsadm -A -t 192.168.2.214:443 -s wlc
ipvs1b21# ipvsadm -a -t 192.168.2.214:80 -r 192.168.1.14 -i
ipvs1b21# ipvsadm -a -t 192.168.2.214:443 -r 192.168.1.14 -i
ipvs1b21# ipvsadm -a -t 192.168.2.214:80 -r 192.168.2.2 -i
ipvs1b21# ipvsadm -a -t 192.168.2.214:443 -r 192.168.2.2 -i
```

The various options used with *ipvsadm* are:

- *-A* adds a virtual server
- *-a* adds a real server
- *-t* specifies that the virtual server is a *tcp* service
- *-s* is a scheduling method
- *-r* represents the real server
- *-i* is an option which specifies that an IPIP tunnel is to be used. IPVS will generate the IPIP packets on its own

You can view the IPVS configuration using the *-L* option with the *ipvsadm* command. The *-n* option saves some DNS look-up time in case the real servers do not have a DNS or host file entry.

```
ipvs1b11# ipvsadm -Ln
IP Virtual Server version 1.2.1 (size=4096)
Prot LocalAddress:Port Scheduler Flags
-> RemoteAddress:Port      Forward Weight ActiveConn InActConn
TCP 192.168.1.214:https wlc persistent 86400
-> 192.168.1.14:https      Tunnel 1  0      0
TCP 192.168.1.214:www rr
-> 192.168.1.14:www       Tunnel 1  0      0
```

A similar output is produced at the *ipvs1b21*. Please make sure that the router at the Data Centre

IPVS Load Balancing Techniques

VS/NAT: This technique is one of the simplest to set up but could present an extra load on the load balancer, because the load balancer needs to rewrite both the request and response packets. The load balancer needs to also act as a default gateway for all the real servers, which does not allow the real servers to be in a geographically different network.

The packet flow in this technique is as follows:

- The load balancer examines the destination address and port number on all incoming packets from the client(s) and verifies if they match any of the virtual services being served.
- A real server is selected from the available ones according to the scheduling algorithm and the selected packets are added to the hash tables recording the connections.
- The destination address and port numbers on the packets are rewritten to match that of the real server and the packet is forwarded to the real server.
- After processing the request, the real server passes the packets back to the load balancer, which then rewrites the source address and port of the packets to match that of the real service and sends it back to the client.

LVS/DR: DR stands for Direct Routing. This technique utilises MAC spoofing and demands that at least one of the load balancer's NIC and real server's NIC are in the same IP network segment as well as the same physical segment. In this technique, the virtual IP address is shared by the load balancer as well as all the real servers. Each real server has a loop-back alias interface configured with the virtual IP address. This loop-back alias interface must be NOARP so that it does not respond to any ARP requests for the virtual IP. The port number of incoming packets cannot be remapped, so if the virtual server is configured to listen on port 80, then real servers also need to service on port 80. The packet flow in this technique is as follows:

2 (where *rproxy2* is located) is set to allow the spoofed address containing the VIP as the source address to pass through. This is the most common problem people face.



Note: If all the real servers are in the same data centre, then you need not worry about router configuration, but in that case you are better off using LVS/DR, which will give you faster performance.

High availability of load balancers—connection synchronisation in IPVS

The *ipvsadm* command can run in daemon mode either as a master or back-up. The master server synchronises the connection information to the back-up server, thus providing a persistent connection in case the master server needs to be failed over to the back-up server.

The high-speed operation of a load balancer can cause some amount of extra load on the load balancers. So we need to make sure that the connection

- The load balancer receives the packet from the client and changes the MAC address of the data frame to one of the selected real servers and retransmits it on the LAN.
- When the real server receives the packet, it realises that this packet is meant for the address on one of its loopback aliased interfaces.
- The real server processes the request and responds directly to the client.

LVS/TUN: This is the most scalable technique. It allows the real servers to be present in different LANs or WANs because the communication happens with the help of the IP tunnelling protocol. The IP tunnelling allows an IP datagram to be encapsulated inside another IP datagram. This allows IP datagrams destined for one IP address to be wrapped and redirected to a different IP address. Each real server must support the IP tunnelling protocol and have one of its tunnel devices configured with the virtual IP. If the real servers are in a different network than the load balancer, then the routers in their network need to be configured to accept outgoing packets with the source address as the virtual IP.

This router reconfiguration needs to be done because the routers are typically configured to drop such packets as part of the anti-spoofing measures. Like the LVS/DR method, the port number of incoming packets cannot be remapped. The packet flow in this technique is as follows:

- The load balancer receives the packet from the client and encapsulates the packet within an IP datagram, and forwards it to a dynamically selected real server.
- The real server receives the packet, 'de-encapsulates' it and finds the inner packet with a destination IP that matches with the virtual IP configured on one of its tunnel devices.
- The real server processes the request and returns the result directly to the user.

synchronisation does not become an overhead to the overall operations. According to the linuxvirtualserver.org website, at least the following connection information needs to be passed to the back-up from the master, which is around 24 bytes:

<Protocol, CIP:CPort, VIP:VPort, RIP:RPort, Flags, State>

Efficient synchronisation is done using UDP multicast inside the Linux kernel. The master load balancer runs the IPVS syncmaster daemon inside the kernel, passing the connection information with the UDP multicast to the back-up load balancer(s) accepting the UDP multicast packet.

On the primary load balancers in each data centre, run the following code:

```
ipvs1b1# ipvsadm --start-daemon=master --mcast-interface=eth1
ipvs1b2# ipvsadm --start-daemon=master --mcast-interface=eth1
```

On the back-up load balancer in each data centre, run the following:

```
ipvslib12# ipvsadm --start-daemon=backup --mcast-interface=eth1
ipvslib22# ipvsadm --start-daemon=backup --mcast-interface=eth1
```

When you want to stop the daemons, you can just run the command given below:

```
# ipvsadm -stop-daemon
```

After starting the daemon on both master and backup, we can now use Heartbeat to provide high availability in our load balancers. I am not detailing the Heartbeat set-up since a similar set-up was discussed in an article that appeared earlier [*Building A Highly Available Reverse-Proxy Using Heartbeat*, Pg 75, March 2009]. So this has been left as an exercise for readers. The important point here is that when the Heartbeat failover occurs, the IP address failover script sends out ARP requests to inform the nodes on the network that the VIP has been failed over and that they should update their ARP cache.

This completes the configuration. However, it would have been a lot better if we could use just one program to do all of the above, i.e., create the virtual server, monitor the virtual server and provide for automatic failover to the back-up, do the connection synchronisation, etc. There are a few tools available to do this. Some of these are *keepalived*, *ultrammonkey* (uses *ldirectord* and Heartbeat, and provides some add-on features), Piranha (a Red Hat favourite), etc.

 **Note:** If you are planning to use *keepalived*, *ultrammonkey* or Piranha, do not execute any of the *ipvsadm* commands described above, as these applications take care of all *ipvsadm* functionalities. And if you have already executed them, it's better to give all machines a reboot to clear them.

Troubleshooting

To start with, add only one real host. In our scenario, I am assuming that you have chosen *rproxy1* (192.168.1.14) as the real host. The client IP for me was 10.1.1.1.

If you can't see the Web page, then first try ping from the client to the VIP.

```
client# ping 192.168.1.214
```

If ping works then run the following *tcpdump* commands on the various servers:

```
director# tcpdump -ln -i eth1 host 192.168.1.14
```

You should see the IPIP tunnel packets here. But if you see any ICMP error packet report that states that it could not connect, then there is a problem at the real server end. Check the tunnel there and make sure that

the link status of the tunl0 interface is marked as UP.

```
realserver# tcpdump -ln -i eth0 host 192.168.1.13
realserver# tcpdump -ln -i tunl0 host 10.1.1.1
realserver# tcpdump -ln -i eth0 host 10.1.1.1
```

If all seems to work well and you can see packets flowing across, then run the following *traceroute* to the client IP address on your real server. The source IP address is spoofed to be the VIP in the packets. If you cannot see the output in this command, then surely your borderline firewall or router is blocking the packets. A sample output is also shown below:

```
realserver# traceroute -n -s 192.168.1.214 10.1.1.1
traceroute to 10.1.1.1 (10.1.1.1) from 192.168.1.214, 30 hops max, 38 byte
packets
1 192.168.1.1 10.280 ms 2.700 ms 2.625 ms
2 10.1.1.1 7.407 ms IC 2.586 ms IC 5.503 ms IC
```

Try to set this up on a local LAN first before moving the set-up to the data centre scenario. And after moving to the data centre scenario, set up one data centre first. This will make troubleshooting easier.

Moving further on

In this four-part series we have seen the setting-up of various components involved in providing a highly available Web infrastructure. We have also seen how to replicate this set-up in multiple data centres and have attempted to utilise all the available capacity by balancing the traffic as evenly as possible across various components located in different data centres.

This is by no means a perfect architecture and was just an attempt to demonstrate the use of various FLOSS components in running a production infrastructure. I sincerely hope that this series has been useful to you. 

Resources

- LVS home page: www.linuxvirtualserver.org
- Ultra Monkey load balance and HA solutions: www.ultrammonkey.org
- Keepalived home page: www.keepalived.org
- IPVS in LVS: www.linuxvirtualserver.org/software/ipvs.html
- KTCPVS in LVS: www.linuxvirtualserver.org/software/ktcpvs/ktcpvs.html
- LVS/TUN mailing list archives: marc.info/?t=100272238300009&w=2&r=1
- LVS/TUN troubleshooting: www.ssi.bg/~ja/TUN-HOWTO.txt
- Red Hat Cluster Suit configuration and management: www.redhat.net/docs/manuals/csgfs/browse/rh-cs-en/index.html

By: Ajitabh Pandey

The author has more than 13 years of diversified IT industry experience in training, support and consulting. You can learn more about him at ajitabhpdny.info and e-mail him at [ajitabhpdny \[at\] ajitabhpdny \[dot\] info](mailto:ajitabhpdny[at]ajitabhpdny[dot]info)



Tips & Tricks



Console-based calculator

 To build a terminal-based calculator, add the following function to your `~/.bashrc` file:

```
function calc
{
    echo "{$1}"|bc -l;
}
```

Reload Bash profile using the following command:

```
$source .bashrc
```

Now run *calc* from terminal as:

```
$ calc 4+7*9-10/5
```

You will get the following output:

```
<---output truncated---->
998 Apr/24 - 15:01:02 cd -
999 Apr/24 - 15:01:48 vi /etc/bashrc
1000 Apr/24 - 15:02:15 history
1001 Apr/24 - 15:03:06 ls
1002 Apr/24 - 15:03:12 ifconfig
1003 Apr/24 - 15:03:15 ls
1004 Apr/24 - 15:03:19 history
1005 Apr/24 - 15:03:49 vi /etc/bashrc
1006 Apr/24 - 15:03:52 history
1007 Apr/24 - 15:04:35 cd home
1008 Apr/24 - 15:04:37 ls
1009 Apr/24 - 15:04:50 vi httpd.conf
1010 Apr/24 - 15:07:37 vi httpd.conf
1011 Apr/24 - 15:07:48 history
<---output truncated---->
```

—NMohana Sundaram, mohan.linux@yahoo.com



Encrypt your file with Vim

Did you know Vim can help you encrypt your file so that no one can open it without knowing the encryption key? Well, all you need to do is to create a file with the `-x` flag:

`vim -x filename`

The above command will prompt you for an encryption key. Provide the key and remember it, as you will need it to open the file from now on.

—Yakub Pasha, linuz_kernel@yahoo.com



Using Vi's *exrc* file

 As you know the `~/.exrc` file is used for making permanent settings to your Vi editor. If you place the following commands in this file, these commands will be available to all your subsequent Vi sessions:

www.usti.cz/STUDY/PERFORMANT/061-062-063-064-065-066

After adding this line, re-login and execute the history command again. Now you will see the commands along with their time of execution. Here is an example:

1) *:abb* is for abbreviation

For example:

```
:abb etl Elitecore Technologies Limited, Ahmedabad
```

So whenever you type *etl* in your Vi editor and press ENTER, SPACE or TAB, *etl* will be replaced by “Elitecore Technologies Limited, Ahmedabad” I use it to create templates for C and C++ programs, like:

```
:abb CPP #include<iostream>^M using namespace std;^M int
main()^M {^M return0;^M }
```

Here ^M is for a new line. Note that to make sure you do not type ^M, for a new line you have to:

- a) *ctrl+v*
- b) press ENTER

It will display ^M as above but it means a new line. So, now every time I need a C++ template, I just type CPP and press ENTER, SPACE or TAB.

2) *:map* is for mapping some command to some shortcut.

For example, the vi command *:set number* is used to display line numbers in a file for every new line. Let's say we want to set a shortcut for it, say, F2. To map this command to shortcut key, the Vi setting will be:

```
:map #2 :set number^M
```

Same for *:set nonumber*; let's set it to F3:

```
:map #3 :set nonumber^M
```

3) *set tabstop=2* is for setting TAB. I want TAB to be equal to 2 spaces, so I set it to 2.

—**Shiv Premani, shivpremani@gmail.com**

Read a CD volume label

To read the volume label of a CD-ROM from the terminal:

```
$ dd if=/dev/cdwriter bs=1 skip=32808 count=32
LFYCD_MARCH08      32+0 records in
32+0 records out
32 bytes (32 B) copied, 4.02731 seconds, 0.0 kB/s
```

Note: The value, /dev/cdwriter depends on your device file.

—**Suhail CK, suhailck@gmail.com**

Disabling ‘shutdown’ with Ctrl+Alt+Del

If you get annoyed, when you accidentally reboot the system by pressing the three magic keys, comment out the following line in your */etc/inittab* file:

```
ca::ctrlaltdel:/sbin/shutdown -t3 -r now
```

You may alternatively allow specific users with this capability by changing the line as below and adding the specific login ids in ‘/etc/shutdown.allow’,

```
ca::ctrlaltdel:/sbin/shutdown -a -t3 -r now
```

—**Ashish Kushwaha, ashish.knit.cs@yahoo.co.in**

More passwd flags

If you change user account details using the *passwd* command—yes, it can do more than changing just the password. Open the new terminal and enter the following commands:

```
passwd -d [user_name]
```

where *-d* deletes the user's password.

Some other useful flags are:

- *-l* locks the user account
- *-u* unlocks the user's account
- *-?* is to get help

—**Thiyagarajan. R, tr.thiyagarajan@yahoo.co.in**

Copy/paste in an easier way

Linux offers Ctrl+C/Ctrl+V paste just like Windows. But it also has a second way of copy/paste: highlight any text and it is automatically copied into X's “Primary” selection buffer—which doesn't change whatever you might have copied with Ctrl+C—and you can paste it anywhere by clicking the middle mouse button. Isn't Linux cool?

—**Bhuvanesh, bhuvibhuvanesh@gmail.com**



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.

CodeSport



Sandy Mannarswamy

Welcome to another installment of 'CodeSport'. This month we take a quick look at the problem of finding out whether a given binary tree is in fact a binary search tree. We then discuss the problem of finding the maximum and minimum in a binary search tree.

Thanks to all the readers who sent in their comments to the problems we discussed in last month's column. Congratulations to Rohit Agarwalla and Rajeev Kumar for getting the solutions to the problem correct. Last month's takeaway problem was on identifying binary search trees. A binary search tree (BST), or ordered binary tree, is a type of binary tree where the nodes are arranged in order: for each node in the tree, all elements in the left sub-tree of the node are less than or equal to the value of the node, and all the elements in the right sub-tree of the node are greater than the node. Binary search trees are used for element insertion and lookup.

Last month's problem was to determine whether a given binary tree is, in fact, a binary search tree or not. Given a plain binary tree, you have to examine the tree to determine if it meets the requirements of a binary search tree, i.e., for every node, all of the nodes in its left tree must be less than or equal to the node, and all the nodes in its right subtree must be greater than the node. The algorithm should return *true* if the binary tree meets the criteria for a binary search tree. Else it should return *false*.

What is the best possible time complexity of such an algorithm? In order to verify that every node in the tree satisfies the binary search tree property, we need to visit each node once. Hence the best possible time complexity can not be less than $O(N)$ where N is the number of nodes in the binary tree. Now how do we come up with an algorithm that verifies whether a binary tree satisfies the binary search tree property at each node?

We use the BST property that a left child of a node is less than or equal to the node, whereas the right child of a node is greater than the node. If we visit the left child first, then the node and then the right child, then these nodes will be visited in order. Such a walk on the nodes of the tree is known as 'in order tree walk'. If we print the nodes as we visit each node, such a visit will print the nodes in monotonically increasing order. Hence all we need to do to verify that a binary tree is a BST, is to visit the nodes of the tree in order and print the elements. If such a walk results in a monotonically increasing sequence of elements, then the tree is a BST.

We can easily write a recursive algorithm to visit the binary tree in order. The steps of our recursive algorithm

at each node x are as follows:

1. Make sure that node x is not NULL
2. Recursively print the elements in the left subtree of x
3. Print the element at node x
4. Recursively print the elements in the right subtree of x

We assume that each node x contains a pointer to its left child in the field 'left', a pointer to its right child in the field 'right', and the field 'key' gives the element value contained in the node x . We can assume the utility functions $left(x)$, which returns the left child, $right(x)$ which gives the right child, and $key(x)$ which gives the value of the element at node x .

```
inOrderTreeWalk(x)
{
    if ( x is not null)
    {
        inOrderTreeWalk(left(x));
        print the value of key(x);
        inOrderTreeWalk(right(x));
    }
}
```

Given below is the pseudo code for the function *IsThisTreeBST*, which takes the root of the given binary tree as a parameter.

```
bool IsThisTreeBST(x)
{
    get S the sequence of elements returned by InorderTreeWalk(x);

    if (S is a monotonically increasing sequence)
        Return true;
    else
        return false;
}
```

I leave it to the reader to write the code to verify that the sequence of elements returned by the *inOrderTreeWalk* is a monotonically increasing sequence or not.

This month's programming question

In this month's column, we will stay with the binary search trees. Given a BST T , how do we find the minimum and

maximum elements present in the tree? Remember that if 'y' is in left subtree of node x, then key(y) is less than or equal to key(x). If 'y' is in right sub-tree of node x, then key(y) is greater than key(x). Since this property is satisfied at each node, the minimum element will be located at the leftmost node of the tree. The maximum element will be located at the rightmost node of the tree.

Given below is the pseudo-code for functions

FindMinimumInBST and *FindMaximumInBST*. These functions take node 'X' which is the root of the BST.

```
FindMinimumInBST(X)
{
    while (left(X) != NULL)
        X = left(X);
    return X;
}

FindMaximumInBST(X)
{
    while (right(X) != NULL)
        X = right(X);
    return X;
}
```

What is the worst case time complexity of *FindMinimumInBST* and *FindMaximumInBST*? These functions traverse the BST from the root to the leaf. Hence their complexity is of the order of $O(h)$ where 'h' is the height of the BST. Now the interesting question is that, given a binary search tree with N elements, what is the best possible height of the tree and what is the worst-case value of the height of the tree? For a complete binary tree with N elements, the height is $O(\log N)$ and this is the best-case value for the height of the tree. The worst-case value occurs when the tree grows linearly in one direction. The worst case value is $O(N)$. I leave it to the reader to come up with examples for the order of tree construction such that we end up with trees of heights $O(\log N)$ and $O(N)$.

BSTs support many dynamic set operations such as *insert*, *delete* and *find*. They can be used for implementing dictionaries and priority queues. Since we have already seen how to find the maximum and minimum, I leave it to the reader to come up with the code for 'search', 'insert' and 'delete'. For insertion and deletion, you have to ensure that the BST property holds good after the operation. All basic operations take time, proportional to the height of the tree. Since the height of the tree in the best-case is $O(\log N)$, the best-case time for these operations is $O(\log N)$. Since the height of the tree in worst-case is $O(N)$, the worst-case time for insert, search and delete operations is $O(N)$. Hence in the worst case, the binary search tree is no better than a simple linked list for supporting these operations. There are various data structures like AVL trees, which are balanced trees. They guarantee a worst-case time of $O(\log N)$ for 'insert', 'delete' and 'search' operations by restructuring the tree so that the height remains $O(\log N)$. We shall discuss them in next month's column.

We next look at another easy problem in the context of BSTs, that of finding the successor and predecessor of a given node. Given a node X with key value equal to key[X], the successor of X is

the node Y such that key[Y] is the smallest key greater than key[X]. If node X is the maximum node in the binary search tree, then X has no successor. Note that finding the successor and predecessor requires no key comparisons. It requires only a tree traversal checking the tree structure.

If node 'X' has a non-empty right subtree, then the successor of X is the minimum node in X's right subtree. If node X has an empty right subtree, X's successor is the node Y for which X is found to be the predecessor—that is, X is the maximum in the left subtree of node Y. Given below is the pseudo-code for finding the successor of a given node X:

```
struct node FindSuccessor(struct node X)
{
    if (right[X] is != NULL)
    {
        return FindMinimumInBST(right[X]);
    }
    Y = parent[X];
    while ((Y != NULL) && (X == right[Y]))
    {
        X = Y;
        Y = Parent[Y];
    }
    return Y;
```

I leave it to the reader to come up with the code for *FindPredecessor*.

This month's takeaway problem

This month's takeaway problem comes from a reader, Rajeev Kumar. This is related to the inversions that form the basis for analysing the time complexity of sorting algorithms.

Inversion of a permutation is defined as follows:

- Permutation: 6 2 3 1 4 7 8 9 5
- Inversion: 3 1 1 1 4 0 0 0 0

Here, 1 has inversion 3, as there are three elements greater than 1 which are left to 1 in the given permutation, and these are 6, 2, and 3. Similarly, 2 has inversion 1 because there is only one element left to 2 which is greater than 2, that is, 6. The problem is to find inversion of a permutation in an efficient manner. It is easy to come with an algorithm with $O(n^2)$ time complexity, but can we find inversion of a permutation in $O(n \log n)$ time complexity even at the cost of some memory?

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me. Feel free to send your solutions and feedback to me at sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming! 

About the author:

Sandy Mannarswamy. The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.

Creating Rhythmic Noise



Get started with Csound, a powerful musical instrument that's playable from Python.

Over a year ago, when I first tried the OLPC/Sugar software, I was fascinated by the activities in TamTam applications. Had any of them been available in my time, I may even have liked attending school.

The TamTam activity builds on top of the Csound framework to allow children to explore sounds and music. As I explored the Csound tutorial by Michel Gogins, I came across his comment, "Of all the languages I have used, both in my career as a programmer and in my career as an algorithmic composer, Python has been

by far the easiest and most productive language to learn and to use."

Gogins also writes, "Csound must be considered as one of the most powerful musical instruments ever created." Hence, the motivation for this article. We will explore how to create noise and, we hope, you will go on to compose music :-) In the current state of the world, we need even more music.

To know a bit about the people behind the code, go to: www.csounds.com/journal/2006spring/meetTheCsound5Developers.html.

Getting started

You will need to install the following packages: *csound* and *csound-python*.

Csound requires an XML-like configuration file with three sections. Create a minimal file, *tutorial.csd*:

```
<CsoundSynthesizer>
<CsOptions>
</CsOptions>
<CsInstruments>
</CsInstruments>
<CsScore>
</CsScore>
</CsoundSynthesizer>
```

You will be satisfied with the default options. The instruments section is where we define the 'orchestra', which comprises all the instruments that will be used to create our musical masterpiece. These are the same as the contents of *myfile.orc* in the beginner's articles in www.csounds.com/journal/articleIndex.html.

The score section will contain a list of instructions for the instruments in the orchestra. These are the same as the contents of *myfile.sco*, mentioned in the articles mentioned earlier.

Create the same simple instrument as in the beginners' introduction, www.csounds.com/ezine/winter1999/beginner/index.html:

```
<CsInstruments>
sr=44100 ; Sample Rate
kr=22050 ; Control Rate
ksmps=2 ; sr/kr As far as I know this is always the case
nchnls=2 ; 1=mono, 2=stereo, 4=quad
instr 1           ; Instrument 1 begins here
aout oscil 10000, 440, 1 ; An oscillator
        outs aout, aout ; Output the results to a stereo sound file
        endin           ; Instrument 1 ends here
</CsInstruments>
```

The first four lines have header information, which controls the output format. The remaining lines are the definition of the instrument, which is a simple oscillator operating at a frequency of 440Hz and a volume of 10000, about a third of the maximum. (Volume is represented as a 16-bit integer.) The third parameter to the *oscil* command/opcode identifies the waveform table to be used in the score given below. The oscillator is given a variable name *aout*. The same sound is passed to both the channels.

Now, add the score:

```
<CsScore>
f 1 0 16384 10 1 ; table #, start time, the size, generator, parameter
i 1 0 1 ; instrument #, start time, duration
</CsScore>
```

f is the waveform table, which is available at the start of performance with 16,384 samples. The generator value 10 with parameter 1 corresponds to a sine wave in Csound.

The *i* line is an instrument event with the instrument number, the start time and the duration in seconds as the parameters.

You can run the following script, get a wave file and play it:

```
$ csound -Wo tutorial.wav tutorial.cs
$ aplay tutorial.wav
```

Controlling Csound using Python

Write the following code in *tutorial.py*. This is pretty useless as all it does is replace the command line for executing Csound. But have patience.

```
# Import the Csound API extension module.
import csnd

# Create an instance of Csound.
csound = csnd.CppSound()

# Enable Csound to print console messages
csound.setPythonMessageCallback()

# Load the tutorial piece created earlier.
csound.load('tutorial.csd')

# Set the Csound command for off-line rendering.
csound.setCommand('csound -Wo tutorial.wav temp.orc temp.sco')

# Export the .orc and .sco file for performance
csound.exportForPerformance()

# Actually run the performance.
csound.perform()
```

We will now add code to generate the score algorithmically. However, our current instrument is pretty hopeless. So, you will need to refine the instrument in *tutorial.csd* as follows:

```
instr 1
iamp = p4
ifqc = p5
itabl1 = p6
aout oscil iamp, ifqc, itabl1
        outs aout, aout
        endin
```

The new oscillator uses the amplitude, frequency and waveform table reference as parameters. The variables in Csound have a strict format reminiscent of Fortran. Local variables start with the letters 'i', 'k' or 'a'. A variable

starting with the letter 'i' is initialised to a value when the instrument is started, and does not usually change. The letter 'a' indicates an audio rate variable and the letter 'k' indicates a control rate variable.

There are also some special 'p' variables or parameters that send values from the score to the orchestra—p1, p2 and p3 are the instrument's number, start time and their duration. The variables p4, p5, p6, etc, are flexible. They are used here for amplitude, frequency and the waveform table.

Now, add a call to *add_score* to the *tutorial.py* as follows:

```
csound.setCommand('csound -W0 tutorial.wav temp.orc temp.sco')
add_score(csound)
# Export the .orc and .sco file for performance
```

Now, code an *add_score* method:

```
def add_score(csound):
    sarega = [130.8, 146.8, 164.8, 174.6, 195.0, 220.0, 246.9, 261.6]
    for time in range(8):
        csound.addNote(1, time, 1, 8000, sarega[time], 1)
```

Your instrument will play the notes. Let us improve the instrument and put an envelope in each event. Modify the instrument definition in *tutorial.csd* as follows:

```
instr 1
iamp = p4
ifqc = p5
itabl1 = p6
kamp linseg 0, .2, 1, .2, .8, p3-.5, .8, .2, 0
aout oscil iamp, ifqc, itabl1
outs aout*kamp, aout*kamp
endin
```

You can define an envelope using the opcode *linseg*, which represents the starting amplitude followed by pairs of time intervals and the amplitude at the end of the interval.

The control variable *kamp* starts with 0, rises to 1 in .2 seconds, then drops to .8 in the next .2 seconds, retaining that value until .2 seconds before the end. In the last .2 seconds, the value drops from .8 to 0. As you can imagine, an instrument can be programmed to generate pretty complex sounds for each event. You can get an idea of the programming possibilities by playing two frequencies close to each other. Replace the *add_score* method in the *tutorial.py* with:

```
def add_score(csound):
    sarega = [130.8, 146.8, 164.8, 174.6, 195.0, 220.0, 246.9, 261.6]
    for time in range(8):
        csound.addNote(1, time, 1, 8000, sarega[time], 1)
        csound.addNote(1, time, 1, 8000, sarega[time] + 5, 1)
```

Not surprisingly, you should hear beats.

Musicians do not work with frequencies. They work with octaves. So, let us define another instrument that uses a converter opcode to convert a number into a frequency. The whole number represents the octave, and the decimal part the semitone. So, at the instruments section in *tutorial.csd*, add the following:

```
instr 2
iamp = p4
ifqc = cpspch(p5)
itabl1 = p6
kamp linseg 0, .2, 1, .2, .8, p3-.5, .8, .2, 0
asigr oscil iamp, ifqc*.999, itabl1
asigr oscil iamp, ifqc*1.001, itabl1
outs asigr*kamp, asigr*kamp
endin
```

Notice that the definition has a different oscillator definition for the right and left channels. In the score section, include a second waveform table:

```
f2 0 16384 10 1.5 .3333
```

This table is also a sine wave but includes the first and second harmonics with amplitudes that are a half and a third of the primary wave.

Now, modify the *add_score* method in *tutorial.py* to use both instruments:

```
def add_score(csound):
    sarega = [130.8, 146.8, 164.8, 174.6, 195.0, 220.0, 246.9, 261.6]
    pitch = [8.00, 8.02, 8.04, 8.05, 8.07, 8.08, 8.11, 9.00]
    for time in range(8):
        csound.addNote(1, time, 1, 8000, sarega[time], 1)
        csound.addNote(2, time, 2, 8000, pitch[time], 1)
```

The second instrument plays the same tones but at an octave higher. It also uses the second waveform table. You can explore various programming possibilities—like varying amplitudes, varying durations and varying the relative start time of instruments.

As you would have noticed, there is no constraint on the size of the orchestra you can single-handedly create. You can find some simple drum instruments at www.csounds.com/ezine/autumn1999/synthesis/index.html. So go ahead, try them and, maybe, create a synthetic tabla!

Next month, we will explore playing around Soundfonts and Csound to create sounds. 

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com



S.G. Ganesh

Silly Programming Mistakes => Serious Harm!

As programmers, we know that almost any software that we use (or write!) has bugs. What we might not be aware of is that many disasters occur because of silly mistakes.

What can software bugs cost? “Nothing,” I hear someone saying. They can be beneficial and ensure ‘job security’— since the more bugs we put in the software, the more work we get in the future to fix those embedded bugs!

On a more serious note, software bugs can even cost human lives. Many mishaps and disasters have happened in the past because of software bugs; see [1] for a detailed list. For example, during the 1980s, at least six people were killed because of a synchronisation bug in the Therac-25 radiation treatment machine. In 1996, the Ariane 5 rocket exploded shortly after its take-off because of an unhandled overflow exception.

A sobering thought about software bugs is that, though they might occur because of silly or innocuous mistakes, they can cause serious harm.

In 1962, the Mariner-I rocket (meant to explore Venus) veered off track and had to be destroyed. It had a few software bugs and one main problem was traced to the following Fortran statement:

DO 5 K = 1, 3. The “.” should have been a comma. The statement was meant to be a DO loop, as in “DO 5 K = 1, 3”, but while typing the program, it was mistyped as “DO 5 K = 1, 3”.

So, what’s the big deal? In old Fortran, spaces were ignored, so we can have spaces in identifiers (yes, believe me, it’s true). Hence this became a declaration for a variable of the real type DO5K with an initial value of 1.3 instead of a DO loop. So, a rocket worth \$18.5 million was lost because of a typo error!

In 1990, the AT&T long distance telephone network crashed for nine hours because of a software bug. It cost the company millions of dollars. The mistake was the result of a misplaced break statement. The code that was put inside a switch statement looked like the following (from [2]):

```
network code()
{
switch (line) {
case THING1:
    doit1();
```

```
        break;
case THING2:
    if (x == STUFF) {
        do_first_stuff();
        if (y == OTHER_STUFF)
            break;
        do_later_stuff();
    } /* coder meant to break to here... */
    initialize_modes_pointer();
    break;
default:
    processing();
} /* ...but actually broke to here! */
use_modes_pointer();/* leaving the modes_pointer uninitialized */
}
```

As you can see, the programmer has put a “break;” after the *if* condition. He actually wanted to break it outside the *if* condition; but the control gets transferred to outside the (enclosing) switch statement! We all know that it is not possible to use “break” to come outside of an *if* block: this simple mistake resulted in a huge loss to AT&T.

Programmers are usually surprised at how silly mistakes such as the use of the wrong operator symbols, the wrong termination condition for a loop, etc, can lead to serious software problems. True, while most such mistakes will not cause any harm, some minor errors could sometimes lead to major disasters. 

References

- Collection of Software Bugs: www5.in.tum.de/~huckle/bugse.html
- Expert C Programming, Peter van der Linden, Prentice Hall PTR, 1994

About the author:

S.G. Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is “60 Tips on Object Oriented Programming”, published by Tata McGraw-Hill. You can reach him at sgganesh@gmail.com.

Testing RESTful WebServices Made Easy

Introducing the WizTools.org RESTClient, a cross-platform tool to test RESTful WebServices and HTTP communications.

 WizTools.org RESTClient is a Java application to test RESTful WebServices. It is also used for testing POX-WebServices (POX: Plain Old XML) over HTTP and other HTTP communications.

History

In early 2007, while working on a series of integration projects, we were using RESTful WebServices. To test our WebServices, I had started the project WizTools.org RESTClient (rest-client.googlecode.com). RESTful WebServices itself was born out of disillusionment with its more complex predecessor, the WS-* stack. The bigger technology companies, sitting in their ivory towers had designed the backbone of SOA (Service Oriented Architecture) using a series of specifications collectively called the WS-* stack. These include: SOAP, WSDL, UDDI, WS-Security and WS-Policy, among others [en.wikipedia.org/wiki/WS-*]. The hacker community, dissatisfied with the introduced complexity of WS-* stack, named a new architecture based on HTTP: RESTful WebServices. The word RESTful WebServices was coined by Roy Fielding in his PhD thesis, “Architectural Styles and the Design of Network-based Software Architectures” [www.ics.uci.edu/~fielding/pubs/dissertation/top.htm].

RESTful WebServices did not aim at protocol independence. It leverages some lesser-known features



of the HTTP protocol. For example, using normal Web browsers, we can make HTTP GET and POST requests. In addition to the common GET and POST, RESTful WebServices use the following:

- PUT
- DELETE



Note: The HTTP specification also defines other request types like HEAD, OPTIONS and TRACE.

Back in 2007, we did not find decent clients to test these types of HTTP requests. Thus WizTools.org RESTClient was born.

Usage

RESTClient requires Java 6. To run the program, use the following command:

```
$ java -jar restclient-ui-2.3-jar-with-dependencies.jar
```

Features

The initial idea behind RESTClient was to acquire the capability to make different kinds of HTTP requests, including GET, POST, PUT and DELETE. The latest version of RESTClient also supports other features like SSL, support for adding custom HTTP headers and

body, HTTP BASIC and DIGEST authentication.

One prominent feature of RESTClient is the ability to save requests, responses and the response body. This is often required for regression testing and proof-of-failure. These options are available inside the *File* menu. The related file extensions are:

1. *.rcq*—The request format. This is a RESTClient-specific XML format storing the request details.
2. *.rcs*—The response format. This also is an XML.
3. *.rcr*—This is the archive of both request and response XMLs compressed as a zip file.

Jetty Servlet Container is embedded inside RESTClient. A servlet that verbosely prints out the request details is attached to it. To start the server, use the *Tools* menu. The default listening port of this server is 10101. This can be changed during RESTClient start-up using the system property *rc:trace-server-port*.

The most powerful feature of RESTClient is its integrated support for tests. RESTClient has the Groovy programming language embedded (Figure 2). So test classes can be written in Groovy. Test classes are based on JUnit 3.x and tests are attached to each request. For example, a simple test would look like what's shown below:

```
public class SampleClassTest
  extends org.wiztools.restclient.RESTTestCase{

  // Test method names should start with 'test':
  public void testStatus(){
    if(response.getStatusCode() != 200){
      fail("Response status is not 200!");
    }
  }
}
```

As you can observe, the test classes need to extend *org.wiztools.restclient.RESTTestCase*. *RESTTestCase* internally extends *junit.framework.TestCase*. The instance of *RESTTestCase* has two predefined instance variables available: *request* and *response*. These instances have various convenient methods to access the *request* and *response* details.

request is of type *org.wiztools.restclient.RoRequestBean*. Useful methods that may be invoked:

- `org.wiztools.restclient.HTTPVersion getHttpVersion()`
- `java.net.URL getUrl()`
- `java.util.Map<String, String> getHeaders()`
- `org.wiztools.restclient.RoReqEntityBean getBody()`

response is of type *org.wiztools.restclient.RoResponseBean*. Some common methods that can be invoked on this, are:

- `int getStatusCode()`
- `java.lang.String getStatusLine()`
- `java.util.Map<String, String> getHeaders()`
- `java.lang.String getResponseBody()`

Both the lists are not exhaustive. Please refer to API docs for a complete list of methods that may be invoked.

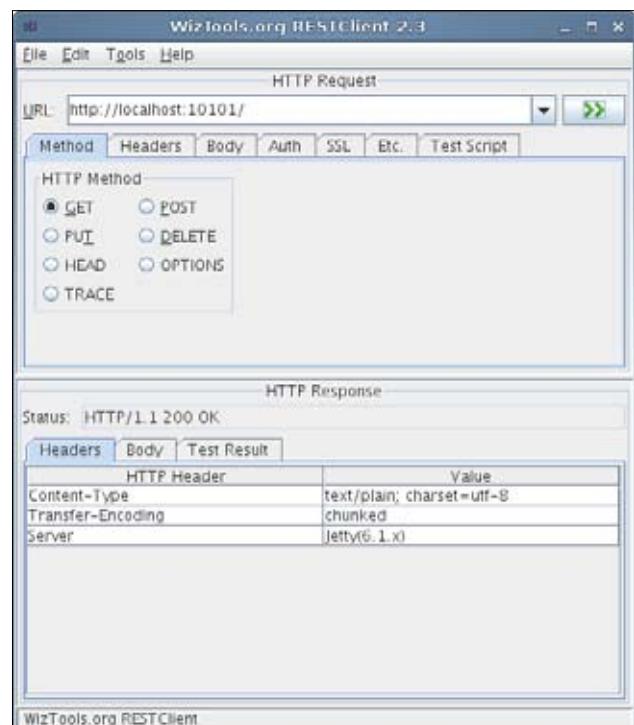


Figure 1: RESTClient Swing interface

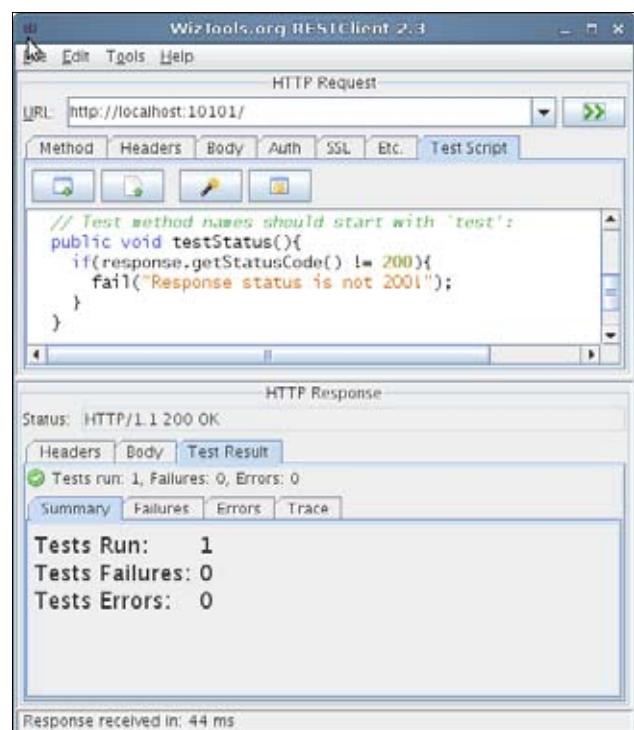


Figure 2: Writing a Groovy test case in RESTClient

The command line

From version 2.3, RESTClient has two binaries: one GUI and one command-line. The command-line tool is used for running requests in a batch and logging their test results. A typical usage is as follows:

```
java -jar restclient-cli-2.3-jar-with-dependencies.jar -o /path/to/responseDir *.rcq
```

This will run all the requests in `*.rcq` files in the current working directory, and save the responses (`*.rcs`) in the `/path/to/responseDir`. The command line client will also print a summary of the test executions.

Extending RESTClient

During the development of version 2.3, the code was re-factored to a more modular form for extensibility. Now the code is organised into various modules managed by Maven. The `restclient-lib` module has the core functionality of RESTClient. Having this as the dependency, various interfaces have been developed (the RESTClient GUI, CLI and Ant plug-in being examples). To demonstrate the ease of the API, I will show you how to write the code to execute a request and write the response-body in the console. First add the dependency for `restclient-lib` in your Maven project [for a detailed discussion on setting up the environment, refer to the Cook Book: code.google.com/p/rest-client/wiki/Cookbook]:

```
<dependency>
<groupId>org.wiztools.restclient</groupId>
<artifactId>restclient-lib</artifactId>
<version>2.3</version>
</dependency>
```

Next, write the following code to execute the request:

```
import org.wiztools.restclient.Request;
import org.wiztools.restclient.RequestBean;
import org.wiztools.restclient.HTTPMethod;
import org.wiztools.restclient.View;
import org.wiztools.restclient.Implementation;
import org.wiztools.restclient.RequestExecuter;

// Step 1: Create the request:
RequestBean requestBean = new RequestBean();
requestBean.setUrl(new java.net.URL("http://wiztools.org/"));
requestBean.setMethod(HTTPMethod.GET);

Request request = requestBean;
```

```
// Step 2: Write the handler
View view = new View(){
    @Override
    public void doStart(Request request){
        // do nothing!
    }

    @Override
    public void doResponse(Response response){
        System.out.println(response.getResponseBody());
    }
}
```

```
@Override
public void doCancelled(){
    // do nothing!
}

@Override
public void doEnd(){
    // do nothing!
}

@Override
public void doError(final String error){
    System.err.println(error);
}

// Step 3: Execute:
RequestExecuter executer = Implementation.of(RequestExecuter.class);
executer.execute(request, view);
```

This example is taken from the *RESTClient Cook Book* [code.google.com/p/rest-client/wiki/Cookbook]. The Cook Book has more details on extending.

What next?

Well, there's Ant and Maven integration. Ant integration work has begun but what about any contributors for Maven?

The team

Various people have contributed to RESTClient through suggestions, ideas, testing efforts, documentation and code contribution. I will not be able to list all of them. But the major contributors, besides me, are:

- Ravi Subramaniam: He was a young and energetic lad. Tragedy struck when we lost him in an accident in 2008. He had contributed the initial persistence code.
- Jacky Chan: No, he is not the actor! Jacky is from China, and contributed various bug fixes and modularisation ideas. But his biggest contribution is in making a RESTClient plug-in for IntelliJ IDEA.
- Velrajan: He has contributed some bug fixes and re-wrote the persistence code using XOM.

Other people who have made substantial contributions are Balasubramani S D and Avi Flax. As I said before, there have been other significant contributors, and I am thankful to all of them. END 

By: Subhash Chandran

Subhash is a software developer working in the innovation department of the Chennai-based software house, Sella Synergy India Ltd. He has contributed to various Open Source projects, and publishes his contributions on his site WizTools.org. He also maintains a technical blog at indiWiz.com. He may be contacted at: subwiz@gmail.com.



Content Management Systems

All you need to make an impressive online presence.

Apart from the ISO images of four FOSS distributions in this month's DVD, we have also managed to pack in some of the best content management systems (CMS). We hope you deploy and test them all. Well, if you really do, let us know your feedback on them, or write a comparison article if you have the time :-)

Drupal is a FOSS modular framework and CMS written in PHP. It is used as a back-end system for many different types of websites, ranging from small personal blogs to large corporate and political sites. The standard release of Drupal, known as "Drupal core", contains basic features common to most CMSs. These include the ability to register and maintain individual user accounts, administration menus, RSS-feeds, customisable layout, flexible account privileges, logging, a blogging system, an Internet forum, and options to create a classic brochure-ware website or an interactive community website.

Joomla CMS enables you to build websites and powerful online applications. Many aspects, including its ease-of-use and extensibility, have made Joomla the most popular website software available. Best of all, Joomla is an open source solution that is freely available to everyone.

WebGUI is a platform for managing all your Web-based content and applications. WebGUI is modular, powerful, secure, and user-friendly. Most users find themselves managing content within hours, and developers can easily plug-in functionality to maximise a site's potential. It is an easy to use content management system, which has ability to create and install custom applications. With WebGUI, you can publish articles, participate in forums, create photo galleries and can even create interactive event calendars.

WordPress is a state-of-the-art Web publishing platform with a focus on aesthetics, Web standards, and usability. It's arguably the de-facto blogging platform.

TYPO3 is a free and open source content management system written in PHP. TYPO3 offers full flexibility and extendibility while featuring an accomplished set of ready-made interfaces, functions and modules. The system is

based on templates. People can choose an existing template and change features such as logo, colours, and fonts, or they can construct their own templates using a configuration language called TypoScript.

Mambo (formerly named Mambo Open Source or MOS) is a free software/open source content management system (CMS) for creating and managing websites through a simple Web interface. It has attracted many users due to its ease of use. Mambo includes advanced features such as page caching to improve performance on busy sites, advanced templating techniques, and a fairly robust API. It can provide RSS feeds and automate many tasks, including web indexing of static pages.

e107 is a content management system written in PHP and using the popular open source MySQL database system for content storage. It's completely free, totally customisable and in constant development.

XOOPS is an extensible, object oriented, easy to use dynamic Web CMS written in PHP. XOOPS is an ideal tool for developing small to large dynamic community websites, intra company portals, corporate portals, blogs and much more.

Plone is a free and open source CMS built on top of the Zope application server. It is suited for an internal website or may be used as a server on the Internet, playing such roles as a document publishing system and group ware collaboration tool. Plone is designed to be extensible.

OpenCms is a professional, easy-to-use website CMS. The fully browser-based user interface features configurable editors for structured content with well-defined fields. Alternatively, content can be created using an integrated WYSIWYG editor similar to well known office applications. A sophisticated template engine enforces a site-wide corporate layout and W3C standard compliance for all content.

Moodle is a Learning Management System (LMS). It is a free Web application that educators can use to create effective online learning sites. It's open source licence and modular design means that people can develop additional functionality. 

Develop Interactive Web Pages with Ease

Here's introducing JQuery, arguably the fastest, most concise and elegant JavaScript framework with support for flexible selectors, CSS3, object detection, method chaining, AJAX, plug-ins and UI effects.

In recent years, many social networking sites and content management systems have cropped up to help people keep in touch with each other. Innumerable applications have been given a Web interface, since users of the applications are comfortable accessing them through a Web browser. With Web 2.0, the areas of interaction design and usability have gone through many developments.

The development of a Web 2.0-enabled application has two aspects—delivering function-rich user interfaces and minimising the efforts spent implementing these common features in applications and websites. The main component of any Web 2.0 application is JavaScript. Till a few years back, JavaScript was considered to be ideal for preliminary functions like client side input validation and form submission. Nowadays, most Web applications get most of their logic-related functions done using JavaScript. Features like validation and auto-suggestion are no longer add-ons—they have become very much part of the basic requirements. There is a paradigm shift, both in user interface technologies and the way JavaScript is used.

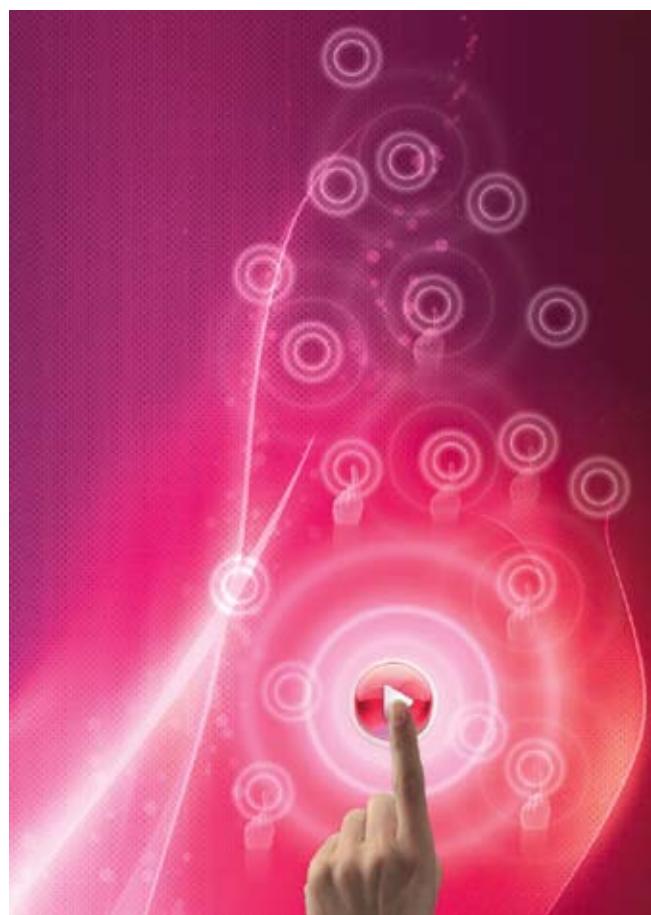
Anyone who has worked on Web application development would have faced enormous issues due to the diversity in standards across the JavaScript engines used in different browsers. We will discuss some of the most challenging aspects:

1. Browser quirks
2. Maintenance of code
3. Tougher implementation of richer client applications

Browser quirks

As we know, not all browsers follow W3C standards. As usual, Microsoft Internet Explorer (IE) likes to live in its

own isolated world, and what's worse is that different versions of IE behave differently. I'm sure you must have developed a killer function using JavaScript, only to find that it completely failed to work in one of the browsers. Things like selection of a value from a select box, event handling, implementing asynchronous calls, etc, are different across browsers. Likewise, having browser-specific code in applications is not really a good idea.



Maintenance of code

Software maintenance is complex and challenging, as it is the longest phase in the software development life cycle. Having discussed differences in JavaScript engines, writing a maintainable code in JavaScript has become even more challenging than an object-oriented language. Easier maintenance of code will help in saving lots of time. Let's discuss a common mistake that we often commit: the structure of the page is mixed with its behaviour. Consider the following code snippet:

```
<html>
  <head>
    <script>
      function showMessage () {
      }
    </script>
  </head>
  <body>
    <div id='details_div' onclick="showMessage()">Click Me</div>
  </body>
</html>
```

In the above case, the presentation (HTML) gets mingled with the behaviour (what gets called when the user clicks the element). We need to separate the presentation from its behavior for better maintainability.

Tougher implementation of rich user interfaces

Though it is possible to create a rich user interface using plain JavaScript, the approach is tougher for beginners and unproductive for development teams. The latter should first work on getting the infrastructure ready and then build their application functionality on top of the infrastructure. But JavaScript libraries help us to write much cleaner code, focusing on solving our problem by giving us the necessary infrastructure. For example, in most JavaScript libraries, features like animation, accessing DOM and special effects have become quite common.

JavaScript libraries

The world is full of problems, and so is software development. Have you ever wondered how so many Web 2.0 applications like WordPress, Orkut and Facebook

ID	Name	Age	Location	Phone
1	Ram	20	Delhi	123566
2	John	21	Chennai	129909234
3	Rahim	30	Mumbai	89345

Figure 1: An example of an HTML table with zebra stripes

were developed? Software development is indebted to the Open Source community for providing world-class libraries, and JavaScript is no exception.

Over the past years, many companies and even individuals have released sophisticated JavaScript libraries. These, by and large, help us to overcome many challenges. Some of the noteworthy JavaScript libraries are GWT, EXTJS, Yahoo, Prototype, DOJO, Open Rico, Mootools and JQuery. Like IDEs such as Eclipse and NetBeans, these libraries have tremendously improved the productivity of developers.

Out of the JavaScript libraries available, JQuery is supposed to be simple, fast and lightweight. This article provides a higher-level overview of JQuery, its features and how to integrate JQuery with Web applications. Towards the end, we will discuss a simple example that covers the basic usage of JQuery. In the coming months, we will also discuss each of the important features of JQuery in detail.

JQuery: An introduction

JQuery is arguably the fastest, most concise and elegant JavaScript framework with support for flexible selectors, CSS3, object detection, method chaining, AJAX, plug-ins and UI effects. It is an open source JavaScript library that simplifies the interaction between HTML and JavaScript. Way back in January 2006, John Resig, the creator of JQuery, announced its first public version.

JQuery is named so to indicate its vibrant ability to 'query' any DOM element from any page. The complete library is written in pure JavaScript. The latest version shows improvements in selection performance compared to older versions. JQuery is designed to change the way we write JavaScript—to write less, and do more. As mentioned earlier, the footprint of JQuery is very small—an uncompressed version of JQuery is 120 KB and a minimised version is about 19 KB in size.

There are innumerable plug-ins available for JQuery.

TechnoMail - Enterprise Email Server
Anti SPAM, Anti Virus, Email Content Filtering

Firewall, Internet Access Control
Content Filtering, Site Blocking

Bandwidth Management System

Managed Email Hosting Solutions

1, Vikas Permes, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

TechnoInfotech[®]

JQueryUI is one plug-in that has a set of components for developing rich client user interfaces. Themeroller [jqueryui.com/themeroller] allows us to design custom JQuery UI themes. Some of the other plug-ins can be found at plugins.jquery.com.

Now, coming back to the subject of this article, it's JQuery and JQuery UI that will help us develop much more interesting and interactive Web pages easily. There's obviously a reason why Google, Digg, Technorati, Mozilla, Wordpress, Drupal and many others use JQuery extensively.

Let's try an example

In order to use JQuery, you need to download it from www.jquery.com. There are two versions of JQuery available: the development and the production version. The production version is compressed and much smaller in size. However, for a trial, any version should do the job. Let us quickly build an HTML table with zebra stripes (which is nothing but a table with alternate rows being presented with different background colours). Apart from zebra stripes, the rows have to be highlighted when the mouse hovers over it. Figure 1 illustrates the example.

For zebra stripes, you need to decide on colours for the table header, odd rows, even rows and the colour you want to use during the mouse-hover. In this example, we have chosen orange for the table header background, white for the odd numbered rows, light blue for even numbered rows and pale green to highlight the rows on mouse hover. The following are the styles used:

```
<style type="text/css">
  th {
    background: orange;
    font-weight: bold;
  }

  td {
    background: white;
  }

  tr.alt td {
    background: palegreen;
  }
</style>
```

By the way, these styles are not a requirement from JQuery alone—you'll need these even when working with any other framework.

The next step is to use JQuery and create the zebra stripes. The following is the code that you need to write to implement the stripes in the table:

```
<script src="jquery-prod.js"></script>
```

```
<script type="text/javascript">
  $(document).ready(function() {
    // JQuery ready is quicker than onload
    $(".Ify tr").mouseover(function() {$(this).addClass("over");});
   mouseout(function()
      {$(this).removeClass("over");});
    $(".Ify tr:even").addClass("alt");
  });
</script>
```

As you can see, you first have to include the JQuery library, which is distributed as a single JavaScript file. We have downloaded it from the JQuery website and renamed it 'jquery-prod.js'. The next script tag contains JavaScript code that uses JQuery functions—*ready*, *mouseover*, *addClass*, *mouseout* and *removeClass*. Also, this is going to work on the HTML tag with the class attribute set to *Ify*.

Let us analyse how the functionalities are achieved. During a page load, JQuery registers the function that needs to be invoked during "mouseover" and "mouseout". For example, when the mouse is moved over "tr", a class attribute "over" is added and when the mouse is moved out, the class attribute is removed. Also, the class attribute of even rows is set to "alt". With these few lines of code, the zebra stripes are ready.

JQuery features

There are many things you need to do to make the Web page interactive with JavaScript. We may need to access a part of the page and modify specific contents, contact the Web server using AJAX and get the data from the server on-demand, change the look and feel of the Web page, manipulate DOM, animate the page, and handle browser events. All this is supported and is relatively easier to do with JQuery. For example, with a few lines of JavaScript, you can implement zebra striped HTML tables with ease. The design and architecture of JQuery is extensible and this is the main reason why JQuery is simple and powerful.

In this article, we looked at some common challenges in using JavaScript and how JavaScript libraries can help overcome them. We also explored how to integrate JQuery with existing Web pages. We hope this article gave you some idea about JQuery. In the coming months, we will discuss more about JQuery and its features. Happy coding! 

By: Lakshmi Narayanan N and Veerabahu Subramanian C

The authors work at HCL Technologies Ltd, Chennai. Lakshmi blogs at grabyourfreedom.blogspot.com and you can contact him at grabyourfreedom@gmail.com. Veerabahu blogs at veechand.wordpress.com and you can contact him at veechand@gmail.com.

EVERYTHING & ANYTHING

32 full active young sports loving men
read Cricket today every month





Subversion 1.6 Security Improvements Illustrated

Subversion 1.6.0 was released on 20th March, 2009 and included a lot of security improvements, which we explore in detail in this article.

Last month, while discussing some of the new and significant features available in Subversion 1.6, we talked a little bit about the security features. This month we'll concentrate exclusively on security improvements.

Warning when passwords are cached in plain text

The Subversion client always had capabilities to cache passwords provided by the user in the `~/.subversion` directory in *NIX systems. Unfortunately, these passwords were cached in plain text form, which the user is not explicitly intimated about. Subversion 1.6 throws a prompt before storing passwords in plain text if it is unable to locate a mechanism to encrypt passwords provided by the operating system. This is illustrated as follows:

```
$ svn co http://localhost/svn/repos wc
Authentication realm: <http://localhost:80> TEST SVN repository
Password for 'stylesen':
-----
ATTENTION! Your password for authentication realm:
```

<http://localhost:80> TEST SVN repository

can only be stored to disk unencrypted! You are advised to configure your system so that Subversion can store passwords encrypted, if possible. See the documentation for details.

You can avoid future appearances of this warning by setting the value of the 'store-plaintext-passwords' option to either 'yes' or 'no' in '`/home/stylesen/.subversion/servers`'.

```
-----
Store password unencrypted (yes/no)? yes
Checked out revision 0.
$
```

The default behaviour is to prompt the user if passwords are going to be stored in plain text. If you don't want to be prompted each time, then you can opt to specify this information in the Subversion servers

configuration file, either globally or on a per-server basis, with the following parameters:

Globally,

```
[global]
store-passwords = yes
store-plaintext-passwords = yes
```

Per server,

```
[groups]
group1 = *.collab.net
othergroup = *.example.com
```

[group1]

```
store-passwords = yes
store-plaintext-passwords = yes
[othergroup]
store-passwords = no
store-plaintext-passwords = yes
```

We already have mechanisms built in Subversion to cache passwords in encrypted form using the `wincrypt` API in Windows and `Keychain` services in Mac OS.

Caching SSL client certificate passphrases

The new version provides a way to cache SSL client certificate passphrases also. Previously, you had an option to specify the client certificate passphrase in the `servers` file, where you needed to hard-code the passphrase for each SSL client certificate with the parameter `ssl-client-cert-pp`. Now Subversion automatically caches this for each certificate. This feature is analogous to caching passwords, where you are prompted before storing the passphrases in plain text, as explained in the above section. You can also have control over the plain text passphrase caching with the following parameters in the Subversion `servers` file:

```
store-ssl-client-cert-pp = (yes/no )
store-ssl-client-cert-pp-plaintext = (yes/no)
```

The following code illustrates caching of plain text passphrases:

```
$ svn co https://localhost/svn/repos wc
Authentication realm: https://localhost:443
Client certificate filename: /home/stylesen/stylesen.p12
Passphrase for '/home/stylesen/stylesen.p12':
-----
ATTENTION! Your passphrase for client certificate:

/home/stylesen/stylesen.p12
```

can only be stored to disk unencrypted! You are advised to configure your system so that Subversion can store passphrase encrypted, if possible. See the documentation for details.

You can avoid future appearances of this warning by setting the value of the 'store-ssl-client-cert-pp-plaintext' option to either 'yes' or 'no' in '/home/stylesen/.subversion/servers'.

```
-----
Store passphrase unencrypted (yes/no)? yes
Checked out revision 0.
```

Encrypted password/passphrase caching

Version 1.6 provides an easy way in which passwords/passphrases can be cached in an encrypted form in *NIX systems. If you are a GNOME user and want to store Subversion passwords in encrypted form, then you can use the GNOME Keyring to do so. KDE users can make use of KWallet. In order to use either of these, the Subversion binaries must be compiled with the option for GNOME Keyring or KWallet support. The password stores that must be used can be configured by specifying it in Subversion config at `~/.subversion/config`.

```
[auth]
### Set password stores used by Subversion. They should be
### delimited by spaces or commas. The order of values determines
### the order in which password stores are used.
### Valid password stores:
### gnome-keyring (Unix-like systems)
### kwallet (Unix-like systems)
### keychain (Mac OS X)
### windows-cryptoapi (Windows)
password-stores = gnome-keyring , kwallet
```

In order to enable Subversion to cache passwords in GNOME Keyring we need to pass the following parameter to the `configure` script while compiling Subversion source:

```
--with-gnome-keyring
```

The above requires GNOME Keyring libraries available in your operating system, failing which Subversion falls back to caching passwords unencrypted. Once you have Subversion binary code

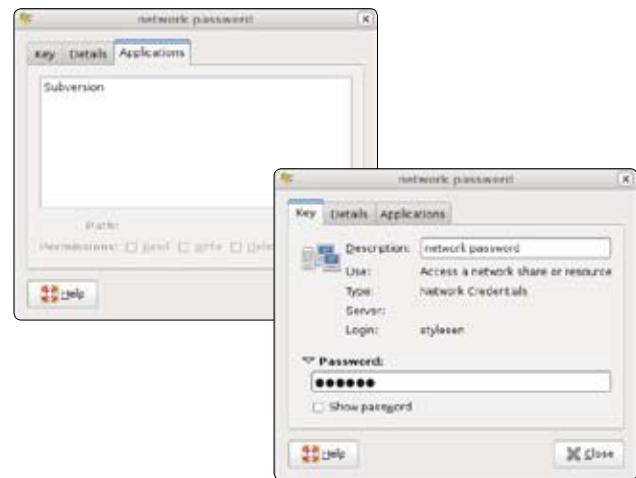


Figure 1: Subversion using GNOME Keyring to store passwords in encrypted form

compiled with GNOME Keyring support, the passwords are automatically cached in Keyring, provided it is unlocked. Figure 1 illustrates a Subversion password stored in GNOME Keyring.

If the GNOME Keyring is locked, you'll get a prompt to unlock it. Once unlocked, the password will get stored in Keyring. A sample run is as follows:

```
$ svn co http://localhost/svn/repos wc
Password for 'default' GNOME keyring:
Authentication realm: <http://localhost:80> TEST SVN repository
Password for 'stylesen':
Checked out revision 0.

$ svn co http://localhost/svn/repos wc
Checked out revision 0.
$
```

KDE can also make use of KWallet in order to store passwords in encrypted form. In order to use KWallet, the Subversion binaries must be compiled with the following option:

```
--with-kwallet
```

Subversion is a widely used version control system in many free software projects and corporate environments. With its new features and improvements, users will benefit from the ease with which they can use it. The latest release (version 1.6.1) of the source can be downloaded from subversion.tigris.org. If you want Subversion binaries for different platforms, visit open.collab.net. 

Senthil Kumaran S.

The author is currently employed by CollabNet at its Version Control group. He is a 'Full Committer' of the Subversion project and a free software enthusiast. To know more about him, visit <http://www.stylesen.org/>

Database Programming in Python

For database programming, the Python DB API is a widely used module that provides a database application programming interface.

rom a construction firm to a stock exchange, every organisation depends on large databases. These are essentially collections of tables, and connected with each other through columns.

These database systems support SQL, the Structured Query Language, which is used to create, access and manipulate the data. SQL is used to access data, and also to create and exploit the relationships between the stored data. Additionally, these databases support database normalisation rules for avoiding redundancy of data.

The Python programming language has powerful features for database programming. Python supports various databases like MySQL, Oracle, Sybase, PostgreSQL, etc. Python also supports Data Definition Language (DDL), Data Manipulation Language (DML) and Data Query Statements.

For database programming, the Python DB API is a widely used module that provides a database application programming interface.

Benefits of Python for database programming

There are many good reasons to use Python for programming database applications:

- Programming in Python is arguably more efficient and faster compared to other languages.
- Python is famous for its portability.
- It is platform independent.
- Python supports SQL cursors.
- In many programming languages, the application developer needs to take care of the open and closed connections of the database, to avoid further exceptions and errors. In

Python, these connections are taken care of. Python supports relational database systems. Python database APIs are compatible with various databases, so it is very easy to migrate and port database application interfaces.

DB-API (SQL-API) for Python

Python DB-API is independent of any database engine, which enables you to write Python scripts to access any database engine. The Python DB API implementation for MySQL is MySQLdb. For PostgreSQL, it supports psycopg, PyGresQL and pyPgSQL modules. DB-API implementations for Oracle are dc_oracle2 and cx_oracle. Pydb2 is the DB-API implementation for DB2. Python's DB-API consists of connection objects, cursor objects, standard exceptions and some other module contents, all of which we will discuss.

Connection objects

Connection objects create a connection with the database and these are further used for different transactions. These connection objects are also used as representatives of the database session.

A connection is created as follows:

```
>>>conn = MySQLdb.connect('library', user='linuxjournal',  
password='python')
```

You can use a connection object for calling methods like *commit()*, *rollback()* and *close()* as shown below:

```
>>>cur = conn.cursor() //creates new cursor object for  
executing SQL statements  
>>>conn.commit() //Commits the transactions  
>>>conn.rollback() //Roll back the transactions
```

```
>>>conn.close() //closes the connection
>>>conn.callproc(proc.param) //call stored procedure for execution
>>>conn.getsource(proc) //fetches stored procedure code
```

Cursor objects

Cursor is one of the powerful features of SQL. These are objects that are responsible for submitting various SQL statements to a database server. There are several cursor classes in *MySQLdb.cursors*:

1. *BaseCursor* is the base class for Cursor objects.
2. *Cursor* is the default cursor class. *CursorWarningMixIn*, *CursorStoreResultMixIn*, *CursorTupleRowsMixIn*, and *BaseCursor* are some components of the cursor class.
3. *CursorStoreResultMixIn* uses the *mysql_store_result()* function to retrieve result sets from the executed query. These result sets are stored at the client side.
4. *CursorUseResultMixIn* uses the *mysql_use_result()* function to retrieve result sets from the executed query. These result sets are stored at the server side.

The following example illustrates the execution of SQL commands using cursor objects. You can use *execute* to execute SQL commands like *SELECT*. To commit all SQL operations you need to close the cursor as *cursor.close()*.

```
>>>cursor.execute('SELECT * FROM books')
>>>cursor.execute("SELECT * FROM books WHERE book_name = 'python' AND
book_author = 'Mark Lutz'")
>>>cursor.close()
```

Error and exception handling in DB-API

Exception handling is very easy in the Python DB-API module. We can place warnings and error handling messages in the programs. Python DB-API has various options to handle this, like *Warning*, *InterfaceError*, *DatabaseError*, *IntegrityError*, *InternalError*, *NotSupportedError*, *OperationalError* and *ProgrammingError*.

Let's take a look at them one by one:

1. *IntegrityError*: Let's look at integrity error in detail. In the following example, we will try to enter duplicate records in the database. It will show an integrity error, *mysql_exceptions.IntegrityError*, as shown below:

```
>>> cursor.execute('insert books values (%s,%s,%s,%s)',('Py9098','Programmin
g With Perl',120,100))
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 95, in
execute
    return self._execute(query, args)
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 114, in
_execute
    self.errorhandler(self, exc, value)
  raise errorclass, errorvalue
_mysql_exceptions.IntegrityError: (1062, "Duplicate entry 'Py9098' for key 1")
```

2. *OperationalError*: If there are any operation errors like no databases selected, Python DB-API will handle this error as *OperationalError*, shown below:

```
>>> cursor.execute('Create database Library')
>>> q='select name from books where cost>=%s order by name'
>>>cursor.execute(q,[50])
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 95, in
execute
    return self._execute(query, args)
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 114, in
_execute
    self.errorhandler(self, exc, value)
  File "/usr/lib/python2.3/site-packages/MySQLdb/connections.py", line 33, in
defaulterrorhandler
    raise errorclass, errorvalue
_mysql_exceptions.OperationalError: (1046, 'No Database Selected')
```

3. *ProgrammingError*: If there are any programming errors like duplicate database creations, Python DB-API will handle this error as *ProgrammingError*, shown below:

```
>>> cursor.execute('Create database Library')
Traceback (most recent call last):>>> cursor.execute('Create database Library')
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 95, in
execute
    return self._execute(query, args)
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 114, in
_execute
    self.errorhandler(self, exc, value)
  File "/usr/lib/python2.3/site-packages/MySQLdb/connections.py", line 33, in
defaulterrorhandler
    raise errorclass, errorvalue
_mysql_exceptions.ProgrammingError: (1007, "Can't create database 'Library'.
Database exists")
```

Python and MySQL

Python and MySQL are a good combination to develop database applications. After starting the MySQL service on Linux, you need to acquire MySQLdb, a Python DB-API for MySQL to perform database operations. You can check whether the MySQLdb module is installed in your system with the following command:

```
>>>import MySQLdb
```

If this command runs successfully, you can now start writing scripts for your database.

To write database applications in Python, there are five steps to follow:

1. Import the SQL interface with the following command:

```
>>> import MySQLdb
```
2. Establish a connection with the database with the following command:

```
>>> conn=MySQLdb.connect(host='localhost',user='root',passwd=')
```

...where *host* is the name of your host machine, followed by the user name and password. In case of the root, there is no need to provide a password.

3. Create a cursor for the connection with the following command:


```
>>> cursor = conn.cursor()
```
4. Execute any SQL query using this cursor as shown below—here the outputs in terms of 1L or 2L show a number of rows affected by this query:


```
>>> cursor.execute('Create database Library')
1L      // 1L Indicates how many rows affected
>>> cursor.execute('use Library')
>>>table='create table books(book_accno char(30) primary key,book_name
      char(50),no_of_copies int(5),price int(5))'
>>> cursor.execute(table)
0L
```
5. Finally, fetch the result set and iterate over this result set. In this step, the user can fetch the result sets as shown below:


```
>>> cursor.execute('select * from books')
2L
>>> cursor.fetchall()
('Py9098', 'Programming With Python', 100L, 50L), ('Py9099', 'Programming
With Python', 100L, 50L)
```

In this example, the *fetchall()* function is used to fetch the result sets.

More SQL operations

We can perform all SQL operations with Python DB-API. Insert, delete, aggregate and update queries can be illustrated as follows.

1. Insert SQL Query

```
>>>cursor.execute('insert books values (%s,%s,%s,%s)',('Py9098','Programmin
g With Python',100,50))
1L      // Rows affected.
>>> cursor.execute('insert books values (%s,%s,%s,%s)',('Py9099','Programmin
g With Python',100,50))
1L  //Rows affected.
```

If the user wants to insert duplicate entries for a book's accession number, the Python DB-API will show an error as it is the primary key. The following example illustrates this:

```
>>> cursor.execute('insert books values (%s,%s,%s,%s)',('Py9099','Programmin
g With Python',100,50))
>>>cursor.execute('insert books values (%s,%s,%s,%s)',('Py9098','Programmin
g With Perl',120,100))
```

Traceback (most recent call last):

```
  File "<stdin>", line 1, in ?
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 95, in
execute
    return self._execute(query, args)
  File "/usr/lib/python2.3/site-packages/MySQLdb/cursors.py", line 114, in
_execute
    self.errorhandler(self, exc, value)
  File "/usr/lib/python2.3/site-packages/MySQLdb/connections.py", line 33, in
defaulterrorhandler
    raise errorclass, errorvalue
_mysql_exceptions.IntegrityError: (1062, "Duplicate entry 'Py9098' for key 1")
```

2. The *Update SQL* query can be used to update existing

records in the database as shown below:

```
>>> cursor.execute('update books set price=%s where no_of_
copies<=%s',[60,101])
2L
>>> cursor.execute('select * from books')
2L
>>> cursor.fetchall()
(('Py9098', 'Programming With Python', 100L, 60L), ('Py9099', 'Programming
With Python', 100L, 60L))
```

3 The *Delete SQL* query can be used to delete existing records in the database as shown below:

```
>>> cursor.execute('delete from books where no_of_copies<=%s',[101])
2L
>>> cursor.execute('select * from books')
0L
>>> cursor.fetchall()
()

>>> cursor.execute('select * from books')
3L
>>> cursor.fetchall() (('Py9099', 'Python-Cookbook', 200L, 90L), ('Py9098',
'Programming With Python', 100L, 50L), ('Py9097', 'Python-Nut shell', 300L,
80L))
```

4. Aggregate functions can be used with Python DB-API in the database as shown below:

```
>>> cursor.execute('select * from books')
4L
>>> cursor.fetchall()
(('Py9099', 'Python-Cookbook', 200L, 90L), ('Py9098', 'Programming With
Python', 100L, 50L), ('Py9097', 'Python-Nut shell', 300L, 80L), ('Py9096', 'Python-
Nut shell', 400L, 90L))
>>> cursor.execute("select sum(price),avg(price) from books where
book_name='Python-Nut shell'")
1L
>>> cursor.fetchall()
((170.0, 85.0))
```

I hope you'll now enjoy the power of Python for database programming! 

References

- The Python home page: www.python.org
- Programming Python by Mark Lutz, O'Reilly Publications
- Accessing Databases using Python DBAPI-2.0 by Federico Di Gregorio: initd.org/pub/software/psycopg/dbapi20programming.pdf

By: Suhas A. Desai

The author works at Tech Mahindra. He writes on open source and security. In his free time, he volunteers for social causes.

**How can mobile solutions
help us increase revenue?**

**How will cloud computing help
us do more with less?**

What technologies can help my business right now?

How can **virtualization** make
our business more efficient?

Will **unified communications**
help us improve customer
responsiveness?



INTEROP®

MUMBAI | 7–9 OCTOBER
Bombay Exhibition Centre

REGISTER TODAY
for a Conference Pass
or Reserve Your
Free Expo Pass.
www.interop.com/mumbai

ANSWERS. ACTION. ADVANTAGE.

The Leading Global Business Technology Event is Coming to India

Get the information you need to build your competitive edge at IT's most comprehensive Conference and Expo. See all the latest innovations—including virtualization, mobility and cloud computing—that will help you increase efficiency, drive revenue and reduce costs.

Register Today for a Conference Pass or Reserve Your Free Expo Pass.
www.interop.com/mumbai

For Exhibiting and Sponsorship Opportunities Contact Pankaj Jain at + 91 98101 72077 or email: pankaj@ubmindia.com.

SPONSORS:

GOULD
FLUKE networks.

BRONZE
Array
Access. Security. Performance.

BADGE
VeriSign

MEDIA SPONSORS:

CRN

EFYtimes.com

ETM

InfoSecurity

InfoStore

iT

LINUX

Network Computing

Toolbox.com

SUPPORTING ASSOCIATIONS:



UBM
United Business Media

techweb

A Voyage to the Kernel



Part 12

Segment 3.1, Day 11



We're entering a new phase in our journey—kernel programming. In the first part, we'll cover a broad introduction to the Linux platform for newbies, along with some history.

The Linux platform

The UNIX platform showed us the power of the multi-layer security architecture, and the beauty of a well-organised structure in terms of subsystems and layers. This is perhaps one aspect that inspired Linus Torvalds, who was then studying computer science at the University of Helsinki, to develop a free operating system. (It should be noted that there is still an ongoing project to remove the non-free portion of the kernel code.)

In 1990 he wrote:

From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)
Newsgroups: comp.os.minix
Subject: Gcc-1.40 and a posix-question
Message-ID: <1991Jul3.100050.9886@klaava.Helsinki.FI>
Date: 3 Jul 91 10:00:50 GMT
Hello netlanders,

Due to a project I'm working on (in minix), I'm interested in the posix standard definition. Could somebody please point me to a (preferably machine-readable format of the latest posix rules? Ftp-sites would be nice.

The 90s home PCs were powerful enough to run a full-blown UNIX OS. So Linus thought it would be a good idea to come out with a freely available academic version of UNIX. He based his project on Minix. To quote him again:

Date: Aug 26 1991, 11:12 am

Subject: What would you like to see most in minix?
To: comp.os.minix

Hello everybody out there using minix -

I'm doing a (free) operating system (just a hobby, won't be big and professional like gnu) for 386(486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhat (same physical layout of the file-system (due to practical reasons) among other things).

I've currently ported bash(1.08) and gcc(1.40), and things seem to work. This implies that I'll get something practical within a few months, and I'd like to know what features most people would want. Any suggestions are welcome, but I won't promise I'll implement them :-)

Linus (torva...@kruuna.helsinki.fi)

PS. Yes - it's free of any minix code, and it has a multi-threaded fs. It is NOT portable (uses 386 task switching etc), and it probably never will support anything other than AT-harddisks, as that's all I have :-(.

He wrote the first Linux kernel in 1991. Two years after Linus' post, there were 12,000 Linux users. Linux is an essentially full UNIX clone. And now the Linux kernel is over six million lines of code! The main reason why Linux gained its popularity is because it was released under the GNU GPL licence (arguably one of the strongest 'copyleft' or quid pro quo licences). Later, Slackware, Red Hat, SuSE, Debian, et al, sprung up to provide packaged

Linux distributions. This further spurred the growth of the GNU/Linux community.

Some of the earlier large-scale users include Amazon, the US Post Office and the German Army. It is interesting to note that Linux machine clusters were used while making movies like *Titanic* and *Shrek*. Today, we have Linux on gadgets like PDAs, mobiles, embedded applications and even on wristwatches. The introduction of 3D acceleration support, support for USB devices, single-click updates of systems and packages made Linux more popular. Desktop users can now log in graphically and start all applications without typing a single character on a terminal. At the same time, you still have the ability to access the core of the system.

Here are a few other utilities and programs (mostly from GNU) that added power to Linux:

- Bash: The GNU shell
- GCC: The GNU C compiler
- coreutils: A set of basic UNIX-style utilities
- findutils: For searching and finding files
- GDB: The GNU debugger
- fontutils: For converting fonts from one format to another or making new fonts
- Emacs: A very powerful editor
- Ghostscript and Ghostview: An interpreter and graphical front-end for the PostScript files
- GNU Photo: A software to manipulate the inputs from digital cameras
- Octave: To perform numerical computations
- GNOME and KDE: GUI desktop environments
- GNU SQL: A relational database system
- Radius: A remote authentication and accounting server

An introduction to the tech side

The Linux kernel is monolithic with module support. Beginners may not fully appreciate the meaning of this statement. A monolithic kernel is essentially a single large complex “do-it-yourself” kernel program, which has several different logical entities, and these run in kernel mode. The Linux kernel has a modular design. When you boot the system, only a minimal resident kernel is loaded into memory. When you request a feature that is not in the kernel, a kernel module (or driver) is dynamically loaded into memory. The object code normally has a collection of functions that implements a file system, a device driver or any other feature at the kernel’s upper layer. This has the following advantages: minimal main memory usage, modularised structure and platform independence.

Our first module

Let’s begin the technical side by writing a module:

```
#include <linux/module.h>
MODULE_AUTHOR("Aasis Vinayak PG");
```

File	Edit	View	Terminal	Tabs	Help
aasisvinayak@GNU-BOX:~\$ lsmod					
Module		Size	Used by		
module		9344	0		
ufs_180859_1		12832	0		
ufs_cp437		13696	0		
vfat		18816	0		
fat		57376	1 vfat		
usb_storage		81728	0		
libusual		27156	1 usb_storage		
ppp_deflate		12888	0		
clib_deflate		28312	1 ppp_deflate		
bsd_comp		13696	0		
ppp_async		17824	0		
crc_ecitt		18112	1 ppp_async		
ppp_generic		32668	3 ppp_deflate, bsd_comp, ppp_async		
slhc		14288	1 ppp_generic		
isofs		49188	1		
udf		88356	0		
crc_itu_t		18112	1 udf		
af_packet		25728	2		
bifmt misc		16984	1		
i915		38144	2		
drm		86856	3 i915		
rfcom		44432	26		

Figure 1: A typical *lsmod* output

```
MODULE_LICENSE("GPL GPLv3");
static int new__init(void)
{
    printk("Let's begin our new segment!\n");
    return 0;
}
static void goodbye__old(void)
{
    printk("Adieu to segment 2!\n");
}
module_init(new__init);
module_exit(goodbye__old);
```

We have written a module. Let’s save this as *module1.c* in the home directory. Now we need a *make* file:

```
obj-m += module.o
all:
    make -C /lib/modules/$(shell uname -r)/build/ M=$(PWD) modules
clean:
    make -C /lib/modules/$(shell uname -r)/build/ M=$(PWD) clean
clean-files := Module.symvers
```

Before proceeding, check the kernel version you are using:

```
aasisvinayak@GNU-BOX:~$ uname -r
2.6.27-7-generic
```

Now we can make our module:

```
aasisvinayak@GNU-BOX:~$ make -C /lib/modules/'uname -r'/build/
M='pwd'
make: Entering directory `/usr/src/linux-headers-2.6.27-7-generic'
  CC [M] /home/aasisvinayak/module1.o
Building modules, stage 2.
MODPOST 1 modules
  CC  /home/aasisvinayak/module1.mod.o
  LD [M] /home/aasisvinayak/module1.ko
make: Leaving directory `/usr/src/linux-headers-2.6.27-7-generic'
```

You may change the directory by changing the *pwd*

command to your directory. We can check our module by issuing the following commands:

```
aasisvinayak@GNU-BOX:~$ sudo dmesg -c > /dev/null
aasisvinayak@GNU-BOX:~$ sudo insmod module1.ko
aasisvinayak@GNU-BOX:~$ sudo dmesg -c
[74091.826941] Let's begin our new segment!
aasisvinayak@GNU-BOX:~$ sudo rmmod module1
aasisvinayak@GNU-BOX:~$ sudo dmesg -c
[74119.021887] Adieu to segment 2!
```

You can see that a *module1.mod.c* file was also created in the process:

```
#include <linux/module.h>
#include <linux/vermagic.h>
#include <linux/compiler.h>

MODULE_INFO(vermagic, VERMAGIC_STRING);

struct module __this_module
__attribute__((section(".gnu.linkonce.this_module")))= {
.name = KBUILD_MODNAME,
.init = init_module,
#ifdefined CONFIG_MODULE_UNLOAD
.exit = cleanup_module,
#endif
.arch = MODULE_ARCH_INIT,
};

static const struct modversion_info __versions[]
__used
__attribute__((section("__versions")))= {
{ 0xa257c5a3, "struct_module" },
{ 0xb72397d5, "printk" },
{ 0xb4390f9a, "mcount" },
};

static const char __module_depends[]
__used
__attribute__((section(".modinfo")))=
"depends=";

MODULE_INFO(srcversion, "4675BE4AD96DEF402B04BD1");
```

You will find that the code requires the following files as well:

```
#include <linux/list.h>
#include <linux/stat.h>
#include <linux/compiler.h>
#include <linux/cache.h>
#include <linux/kmod.h>
#include <linux/elf.h>
#include <linux/stringify.h>
#include <linux/kobject.h>
```

```
#include <linux/moduleparam.h>
#include <linux/marker.h>
#include <asm/local.h>
#include <asm/module.h>
```

If you are an experienced programmer, then by looking at the dependencies you will know how the new file has been created. Novice users need not worry, as this column will address everything from scratch. For the time being, you can just explore the directories */lib/modules/2.x.x-version* and */usr/src/linux-headers-2.x.x-version*.

As I mentioned before, you can load a module by issuing commands. *lsmod* will list the modules (with a small description about its current usage) in your terminal. Figure 1 shows the result of the execution.

You can also view the list of modules to load at boot time by checking the */etc/modules* file. It will have entries similar to the one listed below:

```
# /etc/modules: kernel modules to load at boot time.
#
# This file contains the names of kernel modules that should be loaded
# at boot time, one per line.

fuse
lp
sbp2

# Generated by sensors-detect on Wed Jan 14 21:14:15 2009
# Chip drivers
coretemp
```

You can issue *modprobe* for probing. Its usage is shown below:

```
Usage: modprobe [-v] [-V] [-C config-file] [-n] [-i] [-q] [-Q] [-b] [-o
<modname>] [ -dump-modversions ] <modname> [parameters...]
modprobe -r [-n] [-i] [-v] <modulename> ...
modprobe -l -t <dirname> [ -a <modulename> ...]
```

As this is an introductory column, the kernel build process (which has five parts, viz. Makefile, .config, arch/\$(ARCH)/Makefile, scripts/Makefile, Kbuild Makefiles) will be explained in detail in subsequent articles.

Compiling the kernel

Now we will look at how to compile a kernel:

- Download the latest Linux source. If you have installed Linux with the kernel source option, the source will already be there. You may copy the kernel source to a safe location so that you can restore it later. It is a good idea to backup your */boot* directory as well or at least the kernel image. On my machine, the source was located in */usr/src/linux-2.6.27-7*.
- Create a soft link to your Linux folder, so that you

have */usr/src/linux*:

- `ln -s /usr/src/linux-2.6.27-7 /usr/src/linux`
- Make the necessary changes to the kernel
- `make clean`
- `make mrproper`
- `make clean`
- Configure the modules. You have to specify the modules to be loaded and those to be compiled in the kernel. It would be ideal to start with your existing configuration file.
- `make dep`
- Modify your Makefile EXTRAVERSION tag to give a unique name to your kernel
- Compile the kernel image


```
nohup make bzImage &
tail -f nohup.out
```
- Compile the modules


```
nohup make modules 1> modules.out 2> modules.err &
tail -f modules.err
```
- `make modules_install`
- Copy your kernel image and config file to the */boot* directory:


```
cp arch/i386/boot/bzImage /boot/cpKernel
cp .config /boot/cpKernelConfig
```
- Create a new initrd image so that the modules can be loaded while booting. The module created in the `make modules` step will be located in */lib/modules*:


```
/sbin/mkinitrd /boot/initrd-2.6.27-7cpKernel.img
/lib/modules/2.6.27-7cpKernel
```
- Now you can edit your */boot/grub/menu.lst* file (if you are using the GRUB bootloader), so that while booting you can select which kernel to start. Here is a sample entry in */boot/grub/menu.lst*:


```
timeout 10

splashimage=/boot/grub/splashimages/firework.xpm.gz

title VINU, kernel 2.6.27-7
uuid c21385ba-7c68-4ac0-924f-8bfafadddc5f
kernel /boot/vmlinuz-2.6.27-7
root=UUID=c21385ba-7c68-4ac0-924f-8bfafadddc5f ro quiet splash
initrd /boot/initrd.img-2.6.27-7
quiet

title VINU, kernel 2.6.27-7
uuid c21385ba-7c68-4ac0-924f-8bfafadddc5f
kernel /boot/vmlinuz-2.6.27-7
root=UUID=c21385ba-7c68-4ac0-924f-8bfafadddc5f ro single
initrd /boot/initrd.img-2.6.27-7

title VINU, memtest86+
uuid c21385ba-7c68-4ac0-924f-8bfafadddc5f
kernel /boot/memtest86+.bin
quiet
```
- (Optional step) Delete the entry for the old kernel version in the config file

- Reboot. When the machine restarts, you can select the new kernel!

This is the standard procedure. There are simpler methods as well. (For example, in the case of Debian-based distros, you have a very easy method. Just Google it and find the resources.)

Kernel structure

Within the kernel layer, Linux is composed of five (or six, based on the classification style) major subsystems: the process scheduler (sched), the memory manager (mm), the virtual file system (vfs), the network interface (net), and the inter-process communication (ipc).

Organisation of code

We shall now discuss the way in which the source layout is organised. If you go to the root of the source tree and list the contents, you will see the following directories:

- **arch** – Contains architecture-specific files
- **block** – Has the implementation of I/O scheduling algorithms
- **crypto** – Has the implementation for cipher operations and contains the cryptographic API for implementing encryption algorithms
- **Documentation** – Contains descriptions about various kernel subsystems
- **drivers** – Device drivers for various device classes and peripheral controllers can be found here
- **fs** – Contains the implementation of file systems
- **include** – Kernel header files
- **init** – High-level initialisation and start-up code
- **ipc** – Contains support for Inter-Process Communication (IPC) mechanisms
- **kernel** – Architecture-independent portions of the base kernel
- **lib** – Has the library routines
- **mm** – Holds the memory management implementation
- **net** – Networking protocols
- **scripts** – Scripts used for kernel build
- **security** – Holds the framework for security
- **sound** – Linux audio subsystem
- **usr** – Has the initramfs implementation

We will dedicate a few more sessions to reviewing the kernel structure before we meddle with programming. So look out for the forthcoming columns to hack the kernel! 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language – a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

Post: Lead Analyst
Company: BA Continuum Solutions Pvt. Ltd.
Profile: Incumbent should possess Cognos experience, solid understanding of the Cognos Infrastructure, SQL coding skills, experienced with Teradata or other RDBMS, understanding of WebSphere etc.
Exp.: 9-13
Location: Gurgaon
Email: paramveer.s.raj@bankoffamerica.com

Post: MSS Security Engineering
Company: Symantec Pune
Profile: Should be proficient with Dragon IDS, Checkpoint fw 1/VPN 1, Symantec Enterprise Firewall, Nokia IPSO, Netscreen Firewall, Cisco Pix, Real Secure IDS, SGS Firewall, Manfunt IDS, Windows/2000 configuration and Linux/Unix configuration.
Exp.: 7-10
Location: Chennai
Email: kamlesh_borde@symantec.com

Post: ZOS Development Professional
Company: Objectwin Technology India Pvt. Ltd.
Profile: BS in Engineering, Computer Science, MIS, or related field with 5+ yrs of experience.
Exp.: 5-10
Location: Bangalore
Email: nagendra@objectwin.com

Post: IT Manager
Company: Xilinx India Technology Services Pvt. Ltd.
Profile: Bachelor's degree in information systems/ computer science or equivalent qualification with experience in managing IT operations and capacity planning of IT infrastructure.
Exp.: 10-15
Location: Hyderabad
Email: ravikin@xilinx.com

Post: HP Enterprise Discovery Specialist
Company: Mindlance India Pvt. Ltd.
Profile: Windows/ Intel operating systems proficiency is required. Must possess software installation knowledge on Intel based machines & Unix/Linux. Must have experience in Disaster Recovery procedures and testing schedules.
Exp.: 4-7
Location: Chennai
Email: v.kumar@mndlance.com

Post: Sr. Engineer
Company: Hire At Ease
Profile: Should have experience on networking protocols like TCP/ IP, SNMP.
Exp.: 5-9
Location: Bangalore
Email: jobs@hireatease.com

Post: Peoplesoft Database Administrator
Company: Oceanengineering International Services Ltd.
Profile: Bachelor of Science degree in a related field, with practical experience in administering PeopleSoft 8.8 or higher ERP systems, oracle databases etc.
Exp.: 4-8
Location: Chandigarh
Email: hi.chandigarh@oceaneeringi.com

Post: ASIC Synthesis & STA Sr. Engineer
Company: Naksha Technologies India Pvt. Ltd.
Profile: Expertise in one or more of scripting languages such as Perl/ TCL/ Python/ Shell is a must. Must be proficient of Linux/ Unix based environments.
Exp.: 5-7
Location: Hyderabad
Email: careers@nakshatechnologies.com

Post: Siebel Admin
Company: WOW Global India Pvt. Ltd.
Profile: Should have done Siebel Administration on any of the UNIX flavored operating system.
Exp.: 4-8
Location: Mumbai
Email: resourcing@wowglobal.com

Post: Sr. Oracle DBA
Company: Objectwin Technology India Pvt. Ltd.
Profile: Must be an graduate/ postgraduate in Engineering or MCA with Linux/ Unix experience.
Exp.: 8-10
Location: Chennai
Email: nagendra@objectwin.com

Post: Oracle DBA- Technical Lead
Company: Societe Generale Global Solution Centre
Profile: The candidate should have experience of handling medium to large datawarehouse projects based on Oracle. Should be Oracle certified DBA.
Exp.: 5-8
Location: Bangalore
Email: kaveri.singh@sgcib.com

Post: Technical Support Engineer
Company: PTC Software (Ind) Pvt. Ltd.
Profile: Bachelor's in Mechanical Engineering with experience using Pro/ Engineer. Should have 3D modeling and detailing skills using Pro/ Engineer.
Exp.: 4-7
Location: Pune
Email: ptcindia-wvc@ptc.com

Post: DB2 DBA
Company: Netmagic Solutions Pvt. Ltd.
Profile: Would be responsible for installation of DB, replication/ mirroring, monitoring the DB2, performance tuning of the DB's etc.
Exp.: 4-8
Location: Mumbai
Email: prajakta.narvekar@netmagicsolutions.com

Post: AIX Administrator
Company: Artech Infosystems Pvt. Ltd.
Profile: Should be proficient in OS performance tuning, capacity planning, maintenance patching and new version installation.
Exp.: 5-7, Location: Bangalore
Email: GEMIndia@artechinfo.in

Post: Oracle DBA
Company: NIIT Technologies
Profile: Applicant should be having experience in an Oracle Database/ E-Business Suite Administration, demonstrable experience of operating system control etc. Certification is mandatory.
Exp.: 4-8
Location: Bangalore
Email: RMGBangalore@niit-tech.com

Post: Project Manager- Web Applications
Company: Comspare Infobase Ltd.
Profile: Candidate must be BE/ MCA/ B.Tech and preferably with management background. Should be having knowledge of systems management and project management and technical competence in Database and UNIX/ Linux administration.
Exp.: 8-10
Location: Delhi
Email: sangeeta@infobase.in

Post: Technical Teamleader
Company: Progressive Infotech Pvt. Ltd.
Profile: Should be graduate with good communication skills. Applicant should be able to work under a highly competitive/ pressure situations.
Exp.: 4-9
Location: Mumbai
Email: hrd.mumbai@progressive.in

Post: Embedded Professional- Architect
Company: Xavient Information Services
Profile: Candidate should possess strong documentation, debugging and data analysis skills with ability to multi-task in a fast-paced, deadline driven environment.
Exp.: 5-10
Location: Noida
Email: humanresource@xavient.com

Post: Sr. Software Engineer
Company: Avaya India Pvt. Ltd.
Profile: Bachelor's in Computer science/ ECE with expertise in Linux platform, Unix/ Linux shell Scripting, Perl C, C++, Java Scripting etc. BE/ MS in CS/ ECE with trouble shooting, programming and analytical skills are a pre-requisite.
Exp.: 5-8, Location: Pune
Email: shombe@avaya.com

Post: C++ Engineer
Company: Hire At Ease
Profile: Should be having solid programming skills in C or C++, problem solving, computer algorithms skills with experience in designing and developing a server application.
Exp.: 5-10, Location: Bangalore
Email: jobs@hireatease.com

Post: Oracle DBA
Company: Visnowa Solutions Pvt. Ltd.
Profile: Engineering or MCA with experience in the IT industry as Oracle DBA. Experience on Unix OS such as HP-UX, Solaris, Linux & Scripting is also required.
Exp.: 6-9, Location: Delhi/NCR
Email: rpo@visnowasolutions.com

Post: Sr. Service Delivery Manager
Company: Ritegain
Profile: Candidate should possess strong communication skills in all forms including written, oral, email, telephone and presentation.
Exp.: 7-12, Location: Noida
Email: careers@ritegain.com


Just one right step can help you reach the top
 Log on to naukri.com And take your first step towards success


India's No.1 Job Site

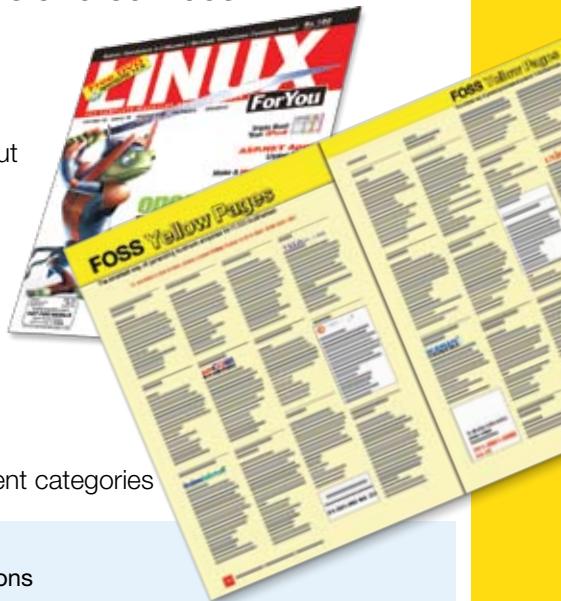
► 2,00,000 Jobs ► 35,000 Recruiters
 ► Get Jobs by email

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues



FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories

TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

 Tear & Send

 Tear & Send

ORDER FORM

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

<input type="checkbox"/> CONSULTANTS	<input type="checkbox"/> HIGH PERFORMANCE COMPUTING	<input type="checkbox"/> SOFTWARE DEVELOPMENT
<input type="checkbox"/> CONSULTANT (FIRM)	<input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS	<input type="checkbox"/> TRAINING FOR PROFESSIONALS
<input type="checkbox"/> EMBEDDED SOLUTIONS	<input type="checkbox"/> LINUX-BASED WEB-HOSTING	<input type="checkbox"/> TRAINING FOR CORPORATE
<input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS	<input type="checkbox"/> MOBILE SOLUTIONS	<input type="checkbox"/> THIN CLIENT SOLUTIONS

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of

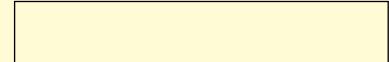
EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card VISA Master Card Please charge Rs. _____

against my credit card No. 

C V V No. _____ (Mandatory)

Date of Birth _____ / _____ / _____ (dd/mm/yy) Card Expiry Date _____ / _____ (mm/yy)



Signature (as on the card)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'llore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Consultant (Firm)

OS3 Infotech

•Silver Solutions Partner for Novell
•High Availability Computing Solutions
•End-to-end Open Source Solutions Provider
•Certified Red Hat Training Partner
•Corporate and Institutional Training
Navi Mumbai
Mobile: 09324113579
Email: info@os3infotech.com
Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX ,OSS & JBOSS related projects. We specialize in high-availability and high-performance clusters, remote and onsite system management, maintenance services, systems planning, Linux & JBOSS consulting & Support services.
Hyderabad
Mobile: 09392493753, Fax: 040-40131726
Email: manojkummar@taashee.com
Web: www.taashee.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek vibraIn B1 UMPC with Linux, Touch Screen, 1 GB RAM 60GB, Wi-Fi, Webcam, upto 6 hour battery (opt.), USB Port, max 1600×1200 resolution, screen 4.8", 7.5"×3.25" Size, weight 526 gm.
New Delhi
Mobile: 09968756177, Fax: 011-26187551
Email: comptekdelhi@compteki.com
Web: www.compteki.com
or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation Education and Training
Mumbai
Tel: 022-28272300, 66462300
Fax: 022-28272399
Email: customercare@aptech.ac.in
Web: www.aptech-education.com,
www.arena-multimedia.com

To advertise in this section, please contact
Somaiah (Bangalore) **09986075717**
Dhiraj (Delhi) **09811206582**

IT-Campus : Academy of Information Technology

IT training and solution company with over 12 years of experience. - RHCE
•Software Training •Hardware Training •Multimedia And Animation •Web Designing •Financial Accounting
Kota (Raj.)
Tel: 0744-2503155, Mobile: 09828503155
Fax: 0744-2505105
Email: m_trilok@yahoo.com
Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization is primarily engaged in Education and Training through its own & Franchise centres in the areas of IT Software, Hardware, Networking, Retail Management and English. The institute also provides customized training for corporates.
New Delhi
Tel: 011-25916832-33
Email: info@mahanindia.com
Web: www.mahanindia.com

Enterprise Communication Solutions

Aware Consultants

We specialize in building and managing Ubuntu/Debian Linux servers and provide good dependable system administration. We install and maintain in-house corporate servers. We also provide dedicated and shared hosting as well as reliable wireless/hybrid networking.
Bangalore
Tel: 080-26724324
Email: sales@aware.co.in
Web: www.aware.co.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty at the Indian Institute of Science, Bangalore and carry over eight decades of experience and fundamental knowledge in the field of DSP and Telecommunication. ESQUBE plays a dominant role in the creation of IP in the domain of Sensors, Signals and Systems.
Bangalore
Tel: 080-23517063
Email: info@esqube.com
Web: www.esqube.com

FOSS Yellow Pages

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open source Enterprise Groupware
Mobile push, Email Syncing of Contacts/Calendar/Tasks with mobiles
•Mail Archival •Mail Auditing •Instant Messaging
New Delhi
Tel: 011-30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

is the training division of Xenitis group of Companies. It is the proud owner of 'Aamar PC', the most popular Desktop brand of Eastern India. These ranges of PC's are sold in the west under the brand name of 'Aamchi PC', in the north as 'Apna PC' and in the south as 'Namma PC'.
Kolkata

Tel: 033-22893280
Email: srinku@xenitisgroup.com
Web: www.techonelabindia.com

netCORE

The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux Based Enterprise Mailing solution with around 1500+ Customer all over India. Key Solutions:
•Enterprise Mailing and Collaboration Solution •Hosted Email Security •Mail Archiving Solution •Push Mail on Mobile •Clustering Solution
Mumbai
Tel: 022-66628000
Mobile: 09322985222
Email: kalpit@netcore.co.in
Web: www.netcore.co.in



Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.
Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis TechnoLab is the first of its kind, state-of-the-art infrastructure, Hardware, Networking and I.T Security training institution headquartered in Kolkata. TechnoLab

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.
New Delhi
Tel: +91-11-26494549
Fax: +91-11-4175 1823
Mobile: +91-9873839960
Email: sales@aisplglobal.com
Web: www.aisplglobal.com

Advent Infotech Pvt Ltd

Advent has an experienced technomarketing team with several years of experience in Networking & Telecom business, and is already making difference in market place. ADVENT qualifies more as Value Added Networking Solution Company, we offers much to customers than just Routers, Switches, VOIP, Network Management Software, Wireless Solutions, Media Conversion, etc.
New Delhi
Tel: 46760000, 09311166412
Fax: 011-46760050
Email: marketing@adventelectronics.com
Web: www.adventelectronics.com

Asset Infotech Ltd

We are an IT solution and training company with an experience of 14 years, we are ISO 9001: 2000. We are partners for RedHat, Microsoft, Oracle and all Major software companies. We expertise in legal software ans solutions.
Dehradun
Tel: 0135-2715965, Mobile: 09412052104
Email: piyush@asset.net.in
Web: www.asset.net.in

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix , Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veritas , Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

FOSS Yellow Pages

HBS System Pvt Ltd

System Integrators & Service Provider. Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symantec Partners of SUN for their new startup E-commerce initiative Solution Provider on REDHAT, SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amitthal@hbsindia.com



Business Open Source

Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

Is the authorised partner for Red-Hat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Srijan Technologies Pvt Ltd

Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of industry verticles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tetrain.com

Web: www.tetrain.com

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

FOSS Yellow Pages

Want to register your organisation in
FOSS Yellow Pages For **FREE***

Call: Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

*Offer for limited period.

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

ManasHosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd

OpenLx Linux with Kalcute (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

Linux Experts

Intaglio Solutions

We are the training and testing partners of RedHat and the first to conduct RHCS exam in delhi for the first time ever.

New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845009939, +91-9343861758

Email : sales@gte-india.com

Web: www.gte-india.com

Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details

Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

DeepRoot Linux Pvt Ltd

DeepRoot Linux is a seven year old GNU/Linux and Free Software company based in Bangalore. We develop Free Software products that are quick-to-deploy and easy-to-use.

Bangalore

Tel: 080-40890000

Email: shivanand@deeproot.in

Web: www.deeproot.in



InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

Noida

Tel: 0102-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: <http://opensource.infoaxon.com>

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Want to register your organisation in
FOSS Yellow Pages*
For **FREE**
CALL: DHIRAJ (DELHI) 09811206582
SOMAJAH (BANGALORE) 09986075717
*Offer for limited period.

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

CATEGORIES FOR FOSS YELLOW PAGES

- Consultants
- Consultant (Firm)
- Embedded Solutions
- Enterprise Communication
- Solutions
- High Performance Computing
- IT Infrastructure Solutions
- Linux-based Web-hosting
- Mobile Solutions
- Software Development
- Training for Professionals
- Training for Corporate
- Thin Client Solutions

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patiala

Mobile: 91-9780531682
Email: pardeep@mrsite.co.in
Web: www.mrsite.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648668, 66091565
Email: ceo@salahsoftware.com
Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing
Mobile: 09880715253
Email: ranga@digitalwaves.in
Web: www.digitalwaves.in

Enjay Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment.

Active Dealer Channel all over India.

Gujarat

Tel.: 0260-3203400, 3241732, 3251732,
Mobile: 09377107650, 09898007650
Email: info@enjayworld.com
Web: www.enjayworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455—66
Fax: 079-27545488
Email: info@bascombridge.com
Web: www.bascombridge.com

Brainnet

Kolkata
Tel: 033-40076450
Email: brainnet@brainware-india.com
Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586
Fax: 022-25706700
Email: contact@cettm.mtnl.in
Web: <http://cettm.mtnl.in/infra>

Complete Open Source Solutions

RHCT, RHCE and RHCSS training.
Hyderabad
Tel: 040-66773365, 9849742065
Email: nayak.sujetu@gmail.com
Web: [www.cossindia.com](http://cossindia.com)

ElectroMech

Redhat Linux and open source solution, RHCE, RHCSS training and exam center, Ahmedabad and Vadodara
Ahmedabad
Tel: 079-40027898
Email: electromech@electromech.info
Web: www.electromech.info



FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2335324
Email: enquiry@focuzinfotech.com
Web: www.focuzinfotech.com

Gujarat Infotech Ltd

GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, MySQL

Ahmedabad

Tel: 079-27452276, Fax: 27414250
Email: info@gujaratinfotech.com
Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

Chennai

Tel: 044-42171278, 9840880558
Email: contactus@lynusacademy.com
Web: www.lynusacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 26600839
Email: info@linuxlearningcentre.com
Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solutions (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking,

while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525
Email: info@mazenesolution.com
Web: www.mazenesolution.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41146565, 32719516
Email: info@netwebindia.com
Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global footprint engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning, Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi

Tel: 011-43612400
Email: info@nhindia.com
Web: www.nhindia.com

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing Corporate and Open classes for RHCE / RHCSS training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

New Delhi

Tel: 46526980-2
Mobile: 09310024503, 09312411592
Email: info@networknuts.net
Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company, Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile: 09873108801
Email: rakhi@stg.in
Web: www.stgonline.com
www.sttglobal.com

TNS Institute of Information Technology Pvt Ltd

Join Red-Hat training and get 100% job guarantee. World's most respected Linux certification. After Red-Hat training, you are ready to join as a Linux Administrator or Network Engineer.

New Delhi

Tel: 011-3085100, Fax: 30851103
Email: nks@tiit.co.in
Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

Kolkata

Tel: 033-22833568, Mobile: 09433111110
Email: enquiry@webelinformatics.com
Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Redhat Training Partner in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses—

RHCSS, SCNA, MCSE, CCNA.

Kolkata

Tel: 033-25488736, Mobile: 09830075018
Email: sinhatuhin1@gmail.com
Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09760099050
Web: www.agamtecindia.com

Amritra Institute of Computer Technology

Amrita Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.

Mobile: 09393733174
Email: aict.hyd@amrita.ac.in
Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicrafts.

Bhopal

Tel: 0755-2661412, 2661559
Fax: 0755-4220022
Email: crisp@crispindia.com
Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced exports & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09312506496
Email: info@cfosdr.com
Web: www.cfosdr.com

CiscoNet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSS, CCNA, MCSE
Kolkata
Tel: 033-25395508, Mobile: 09831705913
Email: info@cisconetinfo.com
Web: www.cisconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state of The Art IT Infrastructure Flexible Batch Timings Available..Leading Networking Institute in Marathwada

Aurangabad

Tel: 0240-3299509, 6621775
Email: aurangabad@cmsinstitute.co.in
Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

Bikaner

Tel: 0151-2202105, Mobile: 09928173269
Email: cmtech.bikaner@gmail.com, kr.gupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer Hardware & Networking

Dehradun

Tel: 3208054, 09897168902
Email: thedisahrainstitute@gmail.com
Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux- Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067566-67, 2609849
Fax: 0172-2615465
Email: info@technoschool.net
Web: <http://technoschool.net>

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer

Hardware Networking, Mobile Phone Maintenance & International Certifications

Jaipur

Tel: 0141-3213378
Email: franchise_gt@gteducation.net
Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

Bhopal

Tel: 0755-4094852
Email: bhopal@hclcdc.in
Web: www.hclcdc.in

IINZTRIX E Technologies Pvt Ltd

No. 1 Training provider in this region.
meerut
Tel: 0121-4020111, 4020222
Mobile: 09927666664
Email: jai@iintrix.com
Web: www.iintrix.com

Indian Institute of Job Oriented Training Centre

Ahmedabad
Tel: 079-40072244—2255—2266
Mobile: 09898749595
Email: info@ijjt.net
Web: www.ijjt.net

Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking
•Software Engg. •Multimedia Training.
Ahmedabad
Tel: 079-32516577, 26607739
Fax: 079-26607739
Email: contact@iantindia.com
Web: www.iantindia.com

IPCC

Bridging Gap with professionals.
Lucknow
Tel: 0522-3919496
Email: ipcciko@yahoo.co.in
Web: www.ipcc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.
New Delhi

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Mobile: 09910710143, Fax: 011-25886909
Email: info@koenig-solutions.com
Web: www.koenig-solutions.com

NACS/CIT
We are Providing Training of LINUX to Professional & Cooperate.
Meerut
Tel: 0121-2420587, Mobile: 9997526668
Email: Info@nacsglobal.com
Web: www.nacsglobal.com

NACS Infosystems (P) Ltd
NACS is a organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.

Meerut
Tel: 0121-2767756, Fax: 0121-4006551
Mobile: 09897796603
Email: info@nacsglobal.com, mohit@nacsglobal.com.
Web: www.nacsglobal.com

Netdiox Computing Systems
We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.

Bangalore
Tel: 080-26640708
Mobile: 09740846885
Email: info@netdiox.com

NetMax-Techologies
Training Partner of RedHat,Cisco
Chandigarh
Tel: 0172-2608315, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com


Netxpert Institute of Advance Networking
Netxpert Noida is a Leading organization to provide Open Source training on RedHat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.
Noida
Tel: 0120-4346847, Mobile: 09268829812
Email: anoop.vyas@netxpriindia.com
Web: www.netxpriindia.com

Netzone Infotech Services Pvt Ltd
Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.

New Delhi
Tel: 011-46015674, Mobile: 9212114211
Email: info@netzoneindia.net

Neuron IT Solutions
We offer end to end services and support to implement and manage your IT Infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration.

Chennai
Mobile: 09790964948
Email: enquiry@neuronit.in
Web: www.neuronit.in

Plexus Software Security Systems Pvt Ltd
Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.

Chennai
Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education
RHCE & RHCSS Certifications
Jabalpur
Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd
Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.

Bangalore
Tel: 080-26639207, 26544135, 22440507
Mobile: +91 9945 282834
Email: counsellors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.

Chandigarh
Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: [stntechologies.com](http://www.stntechologies.com)

South Delhi Computer Centre
SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT) from KARNATAKA STATE OPEN UNIVERSITY.

New Delhi
Tel: 011-26183327, Fax: 011-26143642
Email: southdelhilcomputercentre@gmail.com, southdelhilcomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Systems Quest
Making Tomorrow's professionals TODAY
Bangalore
Tel: 080-41301814
Email: directorv@ssystemsquest.com
Web: www.ssystemsquest.com

Trimax FuturePerfect
A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.
Mumbai
Tel: 022-40681313, Mobile: 09987705638
Fax: 022-40681001
Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd
Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award " REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.

Mumbai
Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses
and authorised testing center of Prometric, Vue and Red Hat.

Mumbai
Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com

Yash Infotech
Authorized Training & Exam Center.
Best Performing Center in Lucknow for RH Training and Examinations.
LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.

Lucknow
Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Want to register your organisation in FOSS Yellow Pages For

FREE

Call: Dhiraj (Delhi) 09811206582

Somaiah (Bangalore) 09986075717

or mail: dhiraj.khare@efyindia.com

or somaiah.km@efyindia.com

*Offer for limited period.

Cool Deal

to Beat the Heat

SUBSCRIBE & Get Assured Gift



24 CD Wallet
(200 nos)
With 1 Year Subscription



Watch with Pen
(500 nos)
With 3 Years Subscription



Executive/Laptop Bag
(500 nos)
With 5 Years Subscription



Magazine	Unit Price (Rs)	5 Years (Rs)	3 Years (Rs)	1 Year (Rs)
Electronics For You (EFY)	40/-	1680 <input type="checkbox"/>	1150 <input type="checkbox"/>	430 <input type="checkbox"/>
Electronics For You (with CD)	60/-	2520 <input type="checkbox"/>	1725 <input type="checkbox"/>	645 <input type="checkbox"/>
Electronics Bazaar	50/-	2100 <input type="checkbox"/>	1440 <input type="checkbox"/>	540 <input type="checkbox"/>
Linux For You (with CD & DVD)	100/-	4200 <input type="checkbox"/>	2880 <input type="checkbox"/>	1080 <input type="checkbox"/>
BenefIT	50/-	2090 <input type="checkbox"/>	1430 <input type="checkbox"/>	530 <input type="checkbox"/>
Facts For You	100/-	4190 <input type="checkbox"/>	2870 <input type="checkbox"/>	1070 <input type="checkbox"/>

- Don't Miss A Single Issue!**
- Ensure Regular Supply**
- FREE Replacement Policy***
- Get Your Money Back if Not Satisfied****
- Order Now For FREE Home Delivery**

* Replacement will be made upon receipt of the damaged / non-receipt copies within 30 days of its publication

** After three months, if you are not satisfied with the magazine, your balance amount will be returned.

Name _____ Company Name _____ Designation _____

Mailing Address _____ Distt. _____

City _____ Pin Code _____ State _____ Phone _____

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓)marked magazine(s), for the period _____ Please find enclosed a sum of Rs _____ by DD/MO/crossed

cheque* bearing the No. _____ dt. _____ in favour of EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs 50 on non-metro cheque)

or charge my credit card VISA MasterCard Please charge Rs. _____

against my credit card No. L L L L L L L L L L L L L L L L

Signature (as on the card)

Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers. # The rates are valid for subscribers within India only. # Please allow 4-6 weeks for processing of your subscription. # Gift against cheque payment will be sent after clearances of the cheque # No discounts will be available for subscription agencies under this scheme. # The models displayed and your gifts received may be slightly different, but will be of the same brand and value. # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

Send this filled-in form or its photocopy to :

Offer Valid till 30th June 2009



EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020; Ph: 011-26810601-03;
Fax: 011-26817563; e-mail: info@efyindia.com; website: www.efyindia.com

Microsoft®

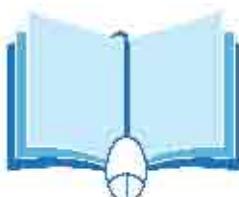
EQUIP YOURSELF

Microsoft invites all developers to join the open source community and online resource centre. Interact with the likes of Sam Ramji, Director of Open Source Software Lab, Jamie Cannon, Lead of OSS Community & Platforms and many others to find solutions to all interoperability challenges.

JOIN WWW.PORT25.IN, A MICROSOFT PLATFORM FOR OPEN SOURCE DEVELOPERS.

PORT25





TECH@EDU EXPO

EDUCATION. POWERED BY TECHNOLOGY.

13-14 June, 09. Pragati Maidan, New Delhi

A TECHNOLOGY SHOW FOR THE EDUCATION INDUSTRY

Demonstrate the state-of-the-art technology to thousands of educational & training Institutes

Mobiles for distance education. Digitisation of curriculum. Innovative teaching aids. Video-conferencing powered classes. State-of-the-art audio & video equipment. Labs that power innovation and R&D...

...pretty exciting stuff, isn't it? But are these technologies really beneficial for educational institutions in India? Have these been implemented by any of the institutions here? If yes, how's been the experience? Most importantly, are these technologies worth the effort and investment?

These are some of the questions that will be addressed at Tech@Edu—a unique event aimed at showcasing the latest technologies that promise to benefit the Indian education sector. You will get to hear technology experts from the leading tech firms. Plus, there will be 'early adopters' from academia who will share their 'gyaan' based on their personal experiences.

Workshops & Seminars

State-of-the-art Labs

Audio & Video Systems

Simpler Management, Thanks to IT Solutions

Hi-tech Teaching Aids

e-Learning: Today & Tomorrow

Harnessing the Web and the Mobile

Open Source Powered Education

Visiting Institutions

- Schools
- Colleges
- Engineering Institutions
- BBA/MBA colleges
- Education Trusts
- Education boards and societies
- Training Institutions
- Examination bodies
- Government bodies

Visitor Profile

- Chairman & Directors
- Senior faculty
- Head of departments
- Professors
- Lecturers
- Teachers
- Purchase Heads
- Trustees
- Principals
- Deans

For speaker slots: Vandana Sharma - 9810286045
email: contact@techedu@efyindia.com

For Booth/Sponsorship: Devendra Kumar - 9810143735
email: sales@techedu@efyindia.com

www.efyedu.com/techedu

An
EFY GROUP
EVENT